

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
17 May 2001 (17.05.2001)

PCT

(10) International Publication Number
WO 01/35209 A2

(51) International Patent Classification⁷: G06F 9/00

(21) International Application Number: PCT/CA00/01339

(22) International Filing Date:
9 November 2000 (09.11.2000)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
09/436,995 9 November 1999 (09.11.1999) US

(71) Applicant (for all designated States except US): UNIVERSITY OF VICTORIA INNOVATION AND DEVELOPMENT CORPORATION [CA/CA]; P.O. Box 3075 STN CSC, R-Hut McKenzie Avenue, Victoria, British Columbia V8W 3W2 (CA).

(72) Inventors; and

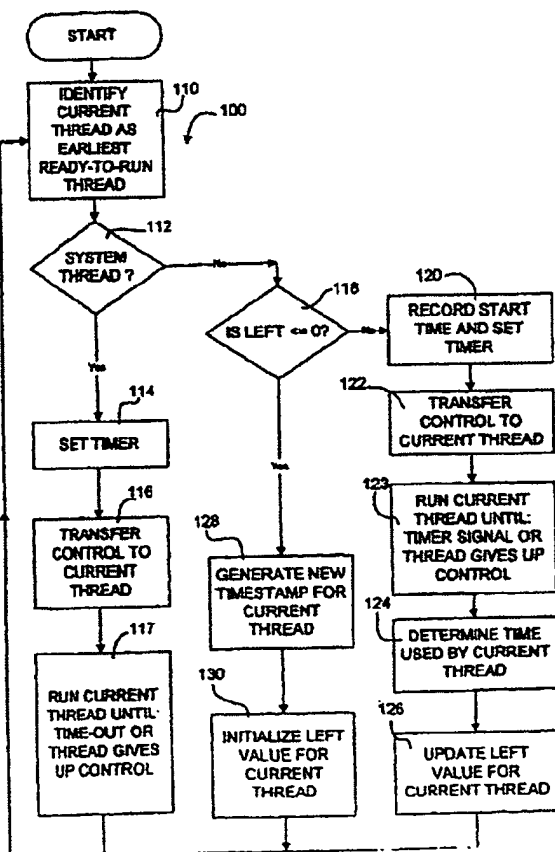
(75) Inventors/Applicants (for US only): PANG, James, C. [CA/CA]; 1703 - 950 Cambie Street, Vancouver, British Columbia V6B 5X5 (CA). SHOJA, Gholamali, C. [CA/CA]; 1650 Barksdale Drive, Victoria, British Columbia V8N 4Z8 (CA). MANNING, Eric, G. [CA/CA]; 2909 Phyllis Street, Victoria, British Columbia V8N 4Z8 (CA).

(74) Agent: MANNING, Gavin, N.; Oyen Wiggs Green & Mutala, 480 - 601 West Cordova Street, Vancouver, British Columbia V6B 1G1 (CA).

(81) Designated States (national): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CR, CU, CZ, DE, DK, DM, DZ, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, US, UZ, VN, YU, ZA, ZW.

[Continued on next page]

(54) Title: MODIFIED MOVE TO REAR LIST SYSTEM AND METHODS FOR THREAD SCHEDULING



(57) Abstract: Methods and systems for scheduling threads in a multi-threaded computer system use a modified move-to-rear list scheduling algorithm. Threads are ordered in a service list according to a virtual time value. System threads always retain a virtual time value. For system threads, the virtual time value serves as a priority. For user threads, the virtual time value is incremented after the thread has received a share of access to the CPU resource. The invention can provide soft real time capability for application software. At the same time, it satisfies system threads which must be executed with some urgency. The thread scheduler of the invention may be used to advantage in the Java multi-threading framework.

WO 01/35209 A2



(84) Designated States (regional): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).

Published:

— Without international search report and to be republished upon receipt of that report.

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

MODIFIED MOVE TO REAR LIST SYSTEM
AND METHODS FOR THREAD SCHEDULING

Technical Field

- 5 This invention relates to scheduling of threads in a multi-threaded computer system. The invention has particular application to managing CPU resources on the JAVA platform.

Background

- 10 Modern general purpose computers make it possible to provide a wide variety of applications which are multimedia in nature. A multimedia application must be run on an underlying supporting environment which can properly support the delivery of continuous media data. Otherwise the multimedia application will not be presented in a satisfactory way. Real-time
- 15 requirements are imposed on the host operating system and its subsystems because continuous media data, such as digital audio or digital video, must be presented continuously at a predetermined rate in order to correctly represent the information being carried. The requirements on the operating system may be classified as being "soft real-time" because, in general, the
- 20 operating system is only required to statistically guarantee that quality of service parameters, such as delay and throughput, will be maintained. Multimedia applications often have deadlines for completing various tasks. Fortunately, missing a particular deadline is generally not fatal as long as the deadline is not missed by too much and as long as most other deadlines
- 25 are not missed.

- Multimedia applications are typically very resource intensive. A supporting computer platform must be able to partition, monitor and police the usage of system resources so that all current application programs can
- 30 make adequate progress, even in the presence of heavy system loads.

- One platform that has several desirable characteristics for multimedia computing is the Java platform. In the Java platform, application software written in the Java programming language runs on top
- 35 of a Java language programming environment. The Java language programming environment includes a Java operating system. The Java operating system runs either on computing hardware which includes a Java processor which directly interprets Java language commands (sometimes known as Java on Java) or on conventional general purpose computing

- 2 -

hardware running an operating system, such as Windows NT or UNIX, and a software Java processor emulator (also known as a Java virtual machine) running on the operating system.

5 Programmers write applications for the Java platform within the Java language programming environment. The Java programming language has a number of advantages. It is small and easily-understood to programmers. Java is object oriented, has built-in support for multi-threaded programming and has automatic memory management. Java is portable,
10 dynamic and robust. In addition an extensive collection of software libraries including Java class libraries and OS specific "native" libraries is available to programmers. The fact that a wide selection of software libraries is available facilitates the rapid development of a wide variety of applications including multimedia processing applications. The interpreted nature of Java
15 enables Java applications to run on a wide variety of computer hardware and a wide variety of operating systems without modification.

 A problem with implementing multimedia applications on Java platforms is that Java presently does not provide support for real-time
20 processing. Java does not provide any mechanism which can be used to monitor, manage or police the usage of the CPU resource to ensure the proper delivery of continuous media data. Current implementations of Java use a static priority-based algorithm for thread scheduling. Such algorithms have been shown to be largely ineffective for multimedia applications.

25 A further problem is that the Java product specification does not specify how thread scheduling should be implemented. Different thread scheduling mechanisms are used by various Java virtual machines. This leads to inconsistencies across platforms and can cause Java applications
30 which function properly on one platform to not function properly on other platforms.

 Nilsen Adding Real Time Capabilities to Java, *Communications of the ACM*, 41(6):49-56, 1998 proposes certain extensions to the Java
35 environment which are designed to facilitate real-time computing. The Nilsen system relies on an off-line analyzer and scheduler to provide

- 3 -

information to a real-time executive for deterministic execution. The Nilsen system is not fully compatible with the Java reference implementation and requires applications to be modified and specifically optimized. The Nilsen system is therefore not useful for general purposes.

5

Sun Microsystems has recently proposed a real-time extension to the Java platform. The proposed extension relies on static priority-based scheduling and exploits real-time facilities in underlying real-time operating systems. Unfortunately, hard real-time scheduling based upon static
10 priorities has been shown to be inefficient for running soft real-time multimedia applications.

Various algorithms suitable for scheduling threads in multi-threaded computer systems are known. Many of these algorithms were
15 developed for scheduling packets in packet switched networks and were subsequently adapted for use in scheduling threads in multi threaded computer systems. For example, weighted fair queuing (WFQ) computes start and finish times for each entity being scheduled and schedules the entities in increasing order of their finish times. The need to compute finish times
20 makes WFQ very computationally intensive. Furthermore, WFQ does not provide fairness when the available bandwidth fluctuates over time due to, for example, sporadic interrupt processing.

Fair queuing based on start time (FQS) is similar to WFQ but
25 schedules entities in increasing order of their start times. FQS is also computationally expensive and does not provide fairness when available bandwidth fluctuates.

Self clocked fair queuing (SCFQ) functions in a similar manner
30 to WFQ but uses an approximation to calculate the finish times for each entity. SCFQ is quite efficient. Unfortunately, SCFQ achieves efficiency at the expense of a maximum scheduling delay which may be unacceptable for many applications, especially multimedia applications.

35 The start time fair queuing (SFQ) algorithm assigns start and finish tags to each entity to be scheduled and then schedules the entities in

increasing order of their start tags. A disadvantage of SFQ is that its delay bound increases linearly with the number of threads in the system. Thus it is not ideal for a general purpose environment in which the number of threads may vary.

5

The experimental Plan 9 operating system incorporates a new scheduling algorithm called Move to Rear List Scheduling (MTR-LS). Move to rear scheduling is described in J. Bruno et al., Move-To-Rear List Scheduling: a new scheduling algorithm for providing QoS guarantees,
10 *Proceedings of the Fifth ACM International Multimedia Conference*, pp. 63-73, November 9-13, 1997. The MTR-LS algorithm assigns a weight, called a service fraction, to each thread. The service fraction specifies the minimum amount of the CPU resource to be allocated to the thread in absolute terms as a fraction of the total available CPU resource. Each thread is also
15 assigned a time stamp. Scheduling of threads is based on their time stamps. The time stamps are adjusted according to the service fractions of their respective threads and the amount of CPU resources consumed by the threads. In addition to the normal QoS guarantees, MTR-LS can provide a
"cumulative service guarantee".

20

A disadvantage of the MTR-LS algorithm is that it cannot effectively be used to schedule system threads which use priorities to specify their urgency. In a priority-based system, the system thread that services
timers and clock interrupts can be given a priority higher than that of any
25 other threads in the system. In such a priority-based system the scheduler can pre-empt the current thread and schedule the timer and clock service thread as soon as the timer and clock service thread becomes runnable. With a move to rear list scheduling algorithm it is not possible to express the relative urgency of different threads since each thread is merely guaranteed
30 a share of the CPU resource but threads are otherwise equal in priority.

There is a need for systems and methods for scheduling threads in multi-threaded computer systems which provide fairness and quality of service guarantees, and yet retain the ability to effectively schedule system
35 threads which must be executed with some urgency. There is a particular need for such systems and methods which can be applied to the Java

operating system because the Java operating system is otherwise well adapted for use in multimedia applications.

Summary of the Invention

5 This invention provides systems and methods which use a modified move to rear list scheduling algorithm. In the modified move to rear list scheduling algorithm each thread is assigned a timestamp value. For high priority system threads the timestamp value is set so that it is always earlier than timestamp values for most other threads. In preferred
10 embodiments the timestamp values for system threads are static. The time stamp values for low priority system threads, such as idler threads are set so that they are always later than the time stamp values of most other threads. The value of the timestamp for the high priority system threads indicates a high priority with which the system threads should be executed
15 if they are ready to run. Conversely, the time stamps of low priority system threads indicate a low priority. All user threads have time stamps that are between those of the high priority system threads and the low priority system threads. For user threads, the timestamp value is reassigned each time the thread uses up an amount of the CPU resource which has been
20 allocated to the thread.

One aspect of the invention provides a computer implemented method for scheduling the running of threads in a multi-threaded computer system. The method comprises maintaining a list of a plurality of threads,
25 each having a timestamp value. A scheduler identifies, as the current thread, the thread which has the earliest timestamp value of any threads which are ready to run. The method maintains a time left value for each thread. If the time left value for the current thread indicates that the thread has some time left to run then control is transferred to the current thread for
30 a running time not exceeding the time left value for the current thread. After the current thread has run for the running time, if the current thread is not a system thread, then the running time is subtracted from the time left value for the thread. Whenever the time left value for the current thread is not greater than zero, the current thread is assigned a new timestamp value
35 which is greater than the timestamp values of any of the other user threads.

- 6 -

The time left value for the current thread is then re-initialized as described below.

5 The methods of the invention permit the fair scheduling of threads while enabling system threads to be scheduled in a manner which takes into account the urgency of the system threads relative to other system threads and user threads.

10 Another aspect of the invention provides a multi-threaded computer system which implements modified move to rear list thread scheduling. The computer system includes: a processor; and a multi-threaded operating system running on the processor. The operating system comprises a plurality of system threads. User software on the computer system comprises one or more user threads. The system includes a memory
15 accessible to the processor. The memory contains a data structure comprising a record for each of a plurality of threads. The plurality of threads comprising at least one high priority system thread and the one or more user threads. Each record includes a timestamp value, a service fraction value and a time left value for the thread corresponding to the record. The computer system
20 includes a scheduler which is adapted to identify as a current thread one of the plurality of threads which is ready to run and has a timestamp value earlier than the timestamp value of any other of the plurality of threads which is ready to run. If the time left value for the current thread is greater than zero, the scheduler transfers control of the processor to the current
25 thread for a running time not exceeding the time left value for the current thread; and, if the current thread is not a system thread, the scheduler subtracts the running time from the time left value for the current thread. If the time left value for the current thread is not greater than zero then the scheduler assigns a new timestamp value to the current thread, the new
30 timestamp value being later than the timestamp values of any of the user threads. The scheduler then initializes the time left value of the current thread and selects a new current thread.

35 Further aspects, advantages and inventive features of the invention are described below.

- 7 -

Brief Description of Drawings

In drawings which illustrate non-limiting embodiments of the invention:

5 Figure 1A is a schematic view of a conventional prior art computing environment;

Figure 1B is a schematic view of a prior art Java computing environment incorporating a Java virtual machine running in a conventional computing environment;

10 Figure 1C is a schematic view of a prior art Java environment including Java software running on a computing hardware which includes a Java processor;

Figure 2 is a schematic view of a computer system according to the invention;

15 Figure 3 is a service list for use in systems and methods according to the invention;

Figure 4 is a flowchart illustrating a method according to the invention;

Figure 5 is a diagram illustrating the assignment of a new timestamp value to a user thread in the method of the invention; and,

20 Figure 6 is a block diagram illustrating functional units of a computer system which includes a modified MTR-LS thread scheduler according to the invention.

Description

25 Figure 1A illustrates a conventional computing environment in which an operating system runs on computer hardware. The computer hardware incorporates one or more processors and provides other resources which may be used by applications. The resources may include memory, access to peripheral devices, and so on. A programming environment
30 mediates interactions between application software and the operating system and interactions between application software and the computer hardware. The application software runs in the programming environment. A problem which has been recognized with the conventional computing environment is that application software must be written specifically for each programming
35 environment since not all programming environments provide the same resources for use by application programs.

- 8 -

Figure 1B shows a prior art Java computing environment. An operating system runs on suitable computing hardware. The operating system may be, for example, Microsoft Windows™ or Sun Solaris™. A Java processor emulator (known as a Java Virtual Machine or "JVM") runs on the operating system. A Java operating system runs on the Java processor emulator. The Java operating system provides resources to a Java language environment. Java application software runs in the Java language environment. Commands which make up the Java application are received by the Java operating system and interpreted by the Java processor emulator. The Java operating system and processor emulator, in turn, cause the operating system and/or hardware to implement the commands. One advantage of the JAVA language environment is that the JAVA language environment is standardized so that every JAVA language environment provides the same resources to JAVA applications without regard to the computing hardware on which the JAVA language environment is provided.

Figure 1C shows an alternative Java computing environment which includes a JAVA processor. A JAVA processor is a microprocessor that is capable of executing directly Java instructions (called "byte codes"). When such a processor is available, the Java operating system and programming environment can be written in the Java language and executed on the processor directly. This eliminates the need for a general purpose operating system and saves much overhead.

The following description explains an embodiment of the invention which provides thread scheduling in a Java computing environment, such as the environment shown in either of Figures 1B or 1C. The invention has particular application to this environment but the methods and systems of the invention may be used for thread scheduling in multi-threaded computer systems generally.

Figure 2 shows a computer system 20 according to the invention. System 20 has a processor 22, a memory 24 accessible to processor 22, and application software 26, which includes a plurality of threads containing instructions to be executed by processor 22. System 20 also includes a software programmable system timer 28 and JAVA operating

- 9 -

system software 30 which includes a scheduler 32 which schedules the running of threads of application software 26.

For each active thread, JAVA operating system software 30
5 maintains a record in a data structure 34 in memory 24. Data structure 34 maintains information about the thread. The form of data structure 34 is not particularly important. Data structure 34 may include information such as the name of the thread; a pointer to stack memory for the thread; a
10 program counter; a register file; machine context information; and a pointer to a parent process as is known in the art. In addition, data structure 34 includes, for each thread, fields for a timestamp value, a service fraction value and a time left value. JAVA operating system 30 also maintains a service list 40, which contains a list of all active threads. Service list 40 is used by scheduler 32 to schedule threads, as described below. In preferred
15 embodiments of the invention, JAVA operating system 30 maintains a queue 42 which contains a list which includes entries for all active threads which are ready to run.

In a computer system 20 according to the invention various
20 types of threads share processor 22. Processor 22 must service high priority system threads. Such threads must be serviced with a high priority for proper operation of computer system 20. An example of a high priority system thread is the thread which services a system clock interrupt and services timeouts generated by timer 28. A second group of threads is the
25 group of threads which make up running user application software 26. Such threads may be called "user threads". The programmers of application software 26 may wish to control how much time processor 22 spends servicing various ones of the user threads. Threads which require large amounts of the processor resource may be given more time to run on
30 processor 22 than other less computationally intensive threads. Finally, computer system 20 may include a number of low priority system threads which perform functions such as running the `finalize()` routines of discarded JAVA objects or performing garbage collection routines.

35 As noted above, data structure 34 includes a timestamp value for each thread. As illustrated in Figure 3, threads in each different group of

- 10 -

threads are assigned timestamp values in a different range. Figure 3 is a map of all possible timestamp values. In the preferred embodiment of the invention, the space of possible timestamp values is subdivided into four ranges, a high priority system thread space 44, a user thread space 46, a low priority system thread space 48, and a reserved space 50. The user thread space 46 is typically much larger than any of the other thread spaces.

High priority system threads are assigned timestamp values from high priority system thread space 44. Since any practical computer system will have a limited number of high priority system threads, high priority system thread space may include a relatively small number of timestamp values. In the preferred embodiment of the invention a clock handler thread is assigned an earliest timestamp of all and a time slicer thread is assigned the next earliest timestamp.

User threads are assigned timestamps in the range 46. Low priority system threads may be assigned timestamps in the area 48. The reserved space 50 is preferably left available for various system tasks, as described below.

The timestamp values assigned to threads are not directly connected to the time as measured by any clock but are numbers. Smaller numbers may be associated with "early" timestamps and larger numbers with "later" timestamps or vice versa. Preferably the timestamp is represented as an integer which permits the timestamp value to have a very large range of values. For example, a 64-bit integer may be used to represent the timestamp value. In general the timestamp value is preferably an integer of 45-bits or more.

The active threads in service list 40 can be conceptually arranged in order of their timestamp values, from earliest to latest. The position of each active thread in such an ordered service list 40 is determined by the timestamp values for the thread. Threads having earlier timestamps (it can be appreciated that "earlier" may be represented by a higher or lower values of the timestamp) appear near the head of service list 40 (i.e. toward high priority system area 44). Threads having later timestamps appear

- 11 -

farther toward the rear of service list 40. Each thread has a service fraction value which indicates how much of the processor's time should be allocated to processing instructions in that thread.

5 It is not necessary to keep the entries in service list 40 ordered by timestamp value since, at any given time, a large number of threads will not be ready to run. At any given time there will be a large number of active threads which are not ready to run. For example, threads may be waiting for a resource, such as a printer, a disk drive, or some other peripheral to become
10 available; threads may be waiting for other threads to complete certain operations or to supply certain results; and/or threads may be waiting for another thread to leave a monitor. Scheduler 32 is only required to schedule threads which are ready to run. Preferably scheduler 32 maintains in queue 42 a list of threads which are ready to run wherein threads with earlier
15 timestamps are closer to the head of queue 42 than are threads which have later timestamps. Queue 42 preferably employs a heap data structure for efficiency.

 As shown in Figure 4, a method 100 for scheduling threads
20 according to the invention begins by identifying a current thread (Step 110). The current thread is the thread in queue 42 which has a timestamp which is earlier than the timestamps of any other threads in queue 42. Where threads in queue 42 are sorted in order of timestamp as described above then scheduler 32 can then be implemented by simply selecting the first thread
25 in the ready to run queue as the current thread.

 System threads are handled differently from user threads. Scheduler 32 determines if the selected thread is a system thread (step 112). If the thread identified as the current thread is a system thread then timer
30 28 is set to measure an interval which determines how long the system thread will be permitted to run (Step 114). Typically the interval is set to equal either a preemption interval or the value of time left specified in data structure 34 for the system thread, whichever is lower. Scheduler 32 then transfers control to the current system thread (Step 116) and the current
35 system thread is allowed to run (step 117) until timer 28 times out or the current system thread voluntarily relinquishes control over processor 22.

- 12 -

Method 100 then returns to Step 110 to identify a new current thread, which may be the same as or different from the thread which has just executed. The time left value for system threads is not altered.

5 Each time a running thread is suspended, context switching code, which is included in scheduler 32, temporarily disables any further signals, calculates the CPU resource consumption for the running thread since the context was switched to it and saves the context of the running thread.

10 If the current thread is not a system thread then a determination is made at Step 118 as to whether or not the time left value for the current thread is greater than zero. If the time left value is greater than zero then the start time of the thread is recorded and timer 28 is set to
15 time out after an interval determined by the shorter of the preemption interval and the time left value for the current thread (Step 120). Once again, timer 28 measures an interval which determines the maximum amount of time that the thread will have to execute without interruption. Each time a user thread begins to run, context switching code records the
20 start time of the thread.

 Method 100 then transfers control to the current user thread (Step 122) and permits the user thread to run until either the thread gives up control or timer 28 times out (Step 123). After the thread has completed
25 execution, the time used by the current thread is determined in Step 124. In Step 126 the time left value for the current thread is then updated by subtracting the time used by the current thread (as determined in Step 124) from the time left value for the current thread.

30 Preferably, if the thread stops running by giving up control before timer 28 times out then at step 126 the time left value for the thread is reduced to zero. As a result, the next time the thread comes to the attention of scheduler 32 it will be moved to the rear of queue 42, as described below, instead of being re-scheduled immediately. A thread may
35 give up control over processor 22 voluntarily when it has finished its current batch of work. A thread may also voluntarily give up control over processor

- 13 -

22 when it is blocked. A thread is blocked when it has further work to do but requires access to some system resource to perform the work. For example, a thread might be blocked when it performs a read from a file and there is no data available, or it performs a monitor entry operation but the monitor is
5 already occupied by another thread.

When a thread is blocked it is no longer "ready to run" and is therefore no longer eligible to participate in the competition among the threads in queue 42 to become the next current thread. Where a thread gives
10 up control of processor 22 because it is finished all of the work that it has to do then it is generally preferable to set the time left value for the thread to zero. If the time left value for the thread is left at a value greater than zero then the thread will still have the earliest time stamp of all user threads, and will be re-scheduled immediately. Resetting the "left" value for a thread
15 which gives up control of processor 22 to zero will ensure that the thread is moved to the rear of queue 42 after all other active user threads in the service list 40. Method 100 then returns to Step 110 to select a new current thread.

20 If, in Step 118, it is determined that the time left value for the current thread is less than or equal to zero then a new timestamp is generated for the current thread in Step 128. Note that Step 128 is never practised on system threads and therefore the timestamp for system threads remains fixed. In Step 130 the time left value for the current thread is
25 initialized. Initializing the time left value for the current thread typically involves multiplying the service fraction for the current thread by a virtual time quantum T . The virtual time quantum T is a user-defined system constant. T is notionally the time which would be taken for all threads in service list 40 to be executed just once in a case where the service fractions
30 for all of the active threads totals 100%.

As illustrated in Figure 5, as each user thread runs on processor 22, its time left value is reduced. When the time left value for the thread has been reduced to zero then the thread is assigned a new timestamp which is
35 later than time stamps of all other user threads in the service list 40. The result is that threads in the user portion 46 of service list 40 leap frog one

- 14 -

another as time passes. In Figure 5, a thread 60 has just completed running. Thread 60 is assigned a new timestamp which is later than the timestamps of any of the four other active user threads in service list 40. Thus, thread 60 is moved to the rear of the list of active service threads, as indicated by arrow
5 62.

Those threads which have a larger service fraction will have a larger time quantum each time their time left values are initialized in step 130. Consequently, threads having a large service fraction are ahead of other
10 user threads in queue 42 will have more time running on processor 22 before being moved behind other user threads in queue 22 than will user threads having a smaller service fraction.

In the preferred embodiment of the invention, timestamps are
15 not re-used. As each thread uses up the quantum of processor time which was specified when its time left value was initialized the thread is assigned a new timestamp which is later than that of any other active user threads. The integer used to represent the timestamp has a large enough range of values that the supply of new timestamps for user threads is practically
20 inexhaustible. For example, the timestamp may be a 64-bit integer value. In this case, if one hundred values of the timestamp were used up every microsecond then it would still take over 5,000 years to use up every possible timestamp value.

25 Low priority system threads, which have timestamps in low priority system area 48 will run only when there are no runnable high priority system threads and no runnable user threads since the timestamp values for low priority system threads are later than the timestamp values for any other threads. Low priority system threads might include, for
30 example, an idler thread, a garbage collection thread, and a finalizer thread which invokes the finalize() function of discarded JAVA objects. The idler thread need not do any useful work. The idler thread is always ready to run. Its presence ensures that there is always at least one runnable thread to be
35 scheduled.

- 15 -

Most preferably the garbage collection thread and the finalizer thread are treated as user threads and are each assigned a small service fraction. If these threads are assigned timestamps in low priority system area 48 which gives them very low priorities then they may not run frequently enough if the system is busy running higher priority threads. This can eventually cause a Java virtual machine to run low on memory. If this happens then the Java virtual machine must pause to run the garbage collector thread. If the garbage collector is run as a user thread in a system according to the invention, on the other hand, then the garbage collector can be guaranteed a small share of CPU time without significantly affecting the operation of user threads.

Timestamps in reserved area 50 may be assigned temporarily to user threads for implementing necessary system maintenance tasks which should not be interrupted by user threads. For example, user threads may be temporarily assigned timestamps in reserved area 50 if a low memory situation develops so that it is necessary to run a garbage collection routine. After memory has been freed by running the garbage collection routine then the previous timestamps of the user threads may be restored.

Figure 6 is a block diagram which shows functional units in a computer system according to one embodiment of the invention. System 200 includes a global priority mapper 210, a scheduler 212, a resource consumption tracker 214, a timer handler 216, and a time slicer 218 each of which may comprise a software thread capable of running on a computer processor. Global priority mapper 210 assigns timestamps to user threads and initializes the time left values for each thread according to the service fraction (CPU resource allocation) for the thread.

Scheduler 212 is invoked each time it is necessary to select a new thread to run. The times when scheduler 212 is invoked to select a new current thread may be called decision epochs. A decision epoch occurs, for example, whenever the current thread is blocked or pre-empted. A decision epoch may occur, for example, when a current thread is pre-empted by a signal handler, when the thread performs some system activity, such as I/O, monitor operations or thread creation, or when it yields control of the

- 16 -

processor. In a Java environment, scheduler 212 may rely on asynchronous signals from an operating system for notification of events such as timer time out or changes in device status. The handler for such a signal will invoke scheduler 212 to cause the currently running thread to be pre-empted and a
5 new thread (which could be the same thread as the currently running thread) to be scheduled. At each decision epoch, scheduler 212 selects a thread for execution. Scheduler 212 scans list 40 (or queue 42) to locate the runnable thread with the highest global priority (earliest timestamp). If the thread identified is not a system thread and has used up its allotment for the
10 current quantum (i.e. has a time left value ≤ 0) then scheduler 212 calls global priority mapper 210 to re-initialize the quantum and global priority (timestamp) for the thread. Global priority mapper 310 resets the priority for each thread to be $\alpha \times T + L$ where α is the service fraction of the thread, T is the virtual time quantum and L is the last time left value for the thread.
15 Scheduler 212 then resumes scanning list 40 (or queue 42) to locate another runnable thread with the highest global priority. Scheduler 212 repeats this operation until it finds a thread with a positive time left value which is ready to run.

20 If the thread identified is a system thread or is a user thread which has not used up all of its allotment for the current quantum then scheduler 212 transfers the control of the CPU to the current thread and sets timer for the time slicer to expire at the end of the next preemption interval, or when the current thread's current quantum is exhausted, whichever
25 comes first.

Resource consumption tracker 214 takes clock readings before the control of the CPU is transferred to a thread and takes clock readings again after the control is given up or revoked. Resource consumption tracker
30 214 then calculates the CPU resource consumption of the thread during the last execution and, if the thread is not a system thread subtracts this amount from the thread's time left value. Resource consumption tracker may be implemented as part of the context switching code which runs whenever context is switched from one thread to another thread.

35

- 17 -

Timer handler 216 services hardware clock interrupts and supports timer time-outs. Timer slicer 218 registers timer timeouts with timer handler 216. When the timer times out then time slicer 218 is invoked (or "woken up") and causes re-scheduling of threads. This prevents any
5 thread from monopolizing the use of the CPU resource. Time slicer 218 interrupts the currently running thread after a time ΔT which is equal to the smaller of the time left value for the thread and a preemption interval. Every time scheduler 212 selects a new thread to run it sets a time-out for time slicer 218 that will expire after ΔT time units. When the time-out
10 expires, time slicer 218 becomes runnable. Because time slicer 218 always has an artificially early timestamp value, the time slicer thread pre-empts the current user thread and causes a re-scheduling. If the current user thread is blocked for some reason, such as a monitor operation, before the time ΔT passes then scheduler 212 cancels the time-out and sets out a new
15 time-out when the next thread is scheduled.

When a new user thread is created the new user thread is preferably assigned a timestamp which is later than that of any other user thread in service list 40. Consequently, even though creating the new thread
20 will lead to a decision epoch, it alone will not cause the current thread to be pre-empted.

In this invention Hoare's "monitor" mechanism (as described in C.A.R. Hoare, Monitors: An Operating System Structuring Concept,
25 *Communications of the ACM*, 17(10):549-557, October, 1974) may be used for thread synchronization. When this is done, signal handling is asynchronous and is split into two parts: a very simple signal handling function is invoked when a signal arrives. This signal handling function invokes a more elaborate signal handling thread. The primary role of the signal handling function is
30 to notify the signal handling thread when a signal is delivered. When a signal is delivered, the current thread is suspended and the signal handling function for the signal is invoked. The signal handling function first invokes context switching code, as described above, to preserve the state of the current thread. The signal handling function then makes its corresponding
35 signal handling thread ready to run. This may be accomplished, for example, by providing a monitor for each signal handling thread and keeping each

- 18 -

signal handling thread in a condition wait queue for its monitor until the signal handling function makes the signal handling thread runnable by setting the condition of the condition wait queue in which the signal handling thread is waiting to true. This removes the signal handling thread from the condition wait queue for its monitor. The signal handling thread becomes
5 runnable immediately. Since the signal handling thread is typically a high priority system thread it will run immediately (unless pre-empted by a higher priority system thread). When the signal handling thread finishes, or is pre-empted, the context switch code is invoked and the scheduler function
10 is called to schedule a new thread. Signal handler functions may use the stack of a pre-empted thread.

As an example, a signal handling function may be associated with a timer 28. The signal handling function is invoked whenever timer 28
15 times out. The signal handling function can cause a clock handler thread to become runnable, and to perform actions which follow from the time-out.

Preemption by Operating System

One problem can occur when it is desired to implement this
20 invention in a multitasking operating system. If a JAVA virtual machine which implements thread scheduling according to the invention is running in such an operating system then the operating system may periodically preempt the JVM and give the processor to another program. If this happens then the computation of the time used by a thread can become inaccurate.
25 Consider, for example, the situation that would occur if a JVM schedules a thread T at time t and the thread runs until a time $t + \Delta t_1$, at which point the operating system takes away control from the JVM (and thread T) until a time $t + \Delta t_1 + \Delta t_2$. Thread T is then allowed to run until time $t + \Delta t_1 + \Delta t_2 + \Delta t_3$. The resource consumption tracker 214 would think that thread T had
30 run for a time $\Delta t_1 + \Delta t_2 + \Delta t_3$ whereas, in reality thread T has only run for $\Delta t_1 + \Delta t_3$ time units.

A solution to this problem is to use the real-time scheduling facility, which is provided by most modern multi-tasking operating systems,
35 to prevent the operating system from arbitrarily preempting the JVM. The JVM may be given a higher priority than any other program running under

- 19 -

the multitasking operating system, including some functions of the operating system itself. To ensure that the operating system has access to the processor for necessary tasks, the JVM can run a thread (an "OS thread") which has a service fraction but does nothing but yield control of the processor to the operating system. For example, 10% of the available processor bandwidth might be allocated to the OS thread.

Timestamp Inversion

Another problem that can occur through the use of the invention is timestamp inversion. Timestamp inversion can occur when monitors are used for synchronization between threads, as is done in Java applications. A monitor is a resource which is available to only one thread at a time. If the monitor is free then a thread can enter the monitor. Otherwise a thread wishing to enter the monitor must wait until the monitor is free. Timestamp inversion can occur when a lower priority thread which is executing inside a monitor blocks a higher priority thread which wishes to enter the same monitor. Consider, for example, the case where T_1 and T_2 are two threads, T_1 is a runnable thread, and T_1 has a timestamp earlier than that of any other runnable thread. T_1 is running inside a monitor M . T_2 has a later timestamp than T_1 and is waiting to enter M . When T_1 finishes its quantum i.e. when its "time left" value becomes zero, scheduler 212 will move thread T_1 to the rear of queue 42 by assigning it a new timestamp. Since T_1 no longer has a timestamp earlier than T_2 , it should no longer block T_2 . However, T_1 will not give up its resources and in particular it will not give up the monitor M it is currently holding. As a result, T_2 cannot proceed, even though it may have an earliest timestamp which is earlier than the timestamp of T_1 which occupies the monitor M , and is blocking T_2 at the entrance of monitor M .

The problem of timestamp inversion may be addressed as follows. When a thread attempts to enter an occupied monitor, the thread is blocked. Each time this occurs, the scheduler compares the timestamp for the thread to the timestamp of the thread which "owns" the monitor. If the thread which is waiting to enter the monitor has an earlier effective timestamp, then the thread which is in the monitor temporarily inherits the earlier timestamp as a new effective timestamp. When the thread in the

- 20 -

monitor exits the monitor, the scheduler restores its former timestamp. Moreover, when a thread which is inside a monitor finishes its quantum (i.e. when its time left value becomes zero) that thread is assigned a new timestamp and its time left value is re-initialized. However, if the monitor's
5 wait queue is not empty (i.e. there are other threads waiting to enter the monitor) then the thread in the monitor would temporarily be given its original timestamp as its effective timestamp until it exits the monitor. When the thread exits the monitor the new timestamp comes into effect. As a result, the thread in the monitor may finish its critical section as quickly
10 as possible and release the monitor for use by other threads. The thread will resume its rightful place in the service list after it leaves the monitor. While this scheme for addressing timestamp inversion problems is simple and allows fair scheduling even when synchronization is required, it may have an adverse effect on some quality of service guarantees, such as any fairness
15 guarantees because the blocking thread is allowed to "borrow" time from its next quantum.

Systems which implement modified move to rear list scheduling according to the invention can provide a cumulative service guarantee. The
20 delay that a particular process experiences will not accumulate over the lifetime of the process and the service rate perceived by the process will not be less than the admitted service rate, as specified by the service fraction, by more than a constant amount. Thus, such systems are well adapted for use by soft real-time applications, such as the delivery of multimedia
25 information.

The invention may be embodied in a program product. The program product comprising any medium which carries a set of computer-readable signals corresponding to instructions which, when run on a
30 computer, cause the computer to execute a method of the invention. The program product may be distributed in any of a wide variety of forms. The program product may comprise, for example, physical media such as floppy diskettes, CD ROMs, DVDs, hard disk drives, flash RAM or the like or transmission-type media such as digital or analog communication links.

35

- 21 -

As will be apparent to those skilled in the art in light of the foregoing disclosure, many alterations and modifications are possible in the practice of this invention without departing from the spirit or scope thereof. Accordingly, the scope of the invention is to be construed in accordance with

5 the substance defined by the following claims.

- 22 -

WHAT IS CLAIMED IS:

1. A computer implemented method for scheduling the running of threads, the method comprising:
 - a) maintaining a list of a plurality of threads, the plurality of threads comprising at least one high priority system thread and at least one user thread and, for each of the plurality of threads, maintaining a timestamp value, a service fraction value and a time left value;
 - b) identifying as a current thread one of the plurality of threads which is ready to run and has a timestamp value earlier than the timestamp value of any other of the plurality of threads which is ready to run;
 - c) if the time left value for the current thread is greater than zero,
 - i) transferring control to the current thread for a running time not exceeding the time left value for the current thread; and,
 - ii) if the current thread is not a system thread, subtracting the running time from the time left value for the current thread; and,
 - d) if the time left value for the current thread is not greater than zero, assigning a new timestamp value to the current thread, the new timestamp value being later than the timestamp values of any of the user threads, and initializing the time left value of the current thread.
2. The method of claim 1 wherein the plurality of threads comprises at least one low priority system thread, the low priority system thread having a timestamp value later than a timestamp value of any of the user threads.
3. The method of claim 1 wherein initializing the time left value of the current thread comprises setting the time left value as the service fraction value of the current thread multiplied by a virtual time quantum.

- 23 -

4. The method of claim 3 wherein the virtual time quantum is a time within which each of the plurality of threads would be executed once if the plurality of threads had service fractions totalling one hundred per-cent.
- 5 5. The method of claim 1 wherein the threads are threads in a JAVA programming environment.
6. The method of claim 5 wherein the high priority system threads
10 comprise a clock handler thread, the clock handler thread servicing a clock interrupt.
7. The method of claim 6 wherein the high priority system threads
15 comprise a time slicer thread, the time slicer thread having a later timestamp than the clock handler thread.
8. The method of claim 5 comprising providing an Operating System which provides priority-based real-time scheduling; running a Java Virtual Machine on the Operating System at a priority higher than a
20 priority allocated to the Operating System; and providing an OS thread running on the Java Virtual Machine, the OS thread providing CPU resources to the Operating System.
9. The method of claim 5 wherein the user threads comprise a garbage
25 collection thread and a finalizer thread wherein the garbage collection thread runs a garbage collection routine and the finalizer thread runs finalize() routines of discarded JAVA objects.
10. The method of claim 5 wherein one of the user threads comprises an
30 OS thread, the OS thread providing CPU resources to an underlying operating system.
11. The method of claim 1 comprising maintaining a queue containing
35 records of all ready-to-run threads in the list, the records arranged in the queue in order of the timestamp values for the corresponding threads.

- 24 -

12. The method of claim 1 wherein the timestamp comprises an integer having in excess of 45 bits.
- 5 13. The method of claim 7 wherein the time slicer thread uses a timer to discontinue execution of a thread after the lesser of a predetermined preemption time and the time left value for the thread.
- 10 14. The method of claim 5, comprising running the threads on a computer comprising a processor capable of directly executing Java byte codes.
- 15 15. The method of claim 5, comprising running the threads on a computer comprising a processor running an operating system, a computer program emulating a Java processor and a Java operating system.
- 15 16. The method of claim 5 wherein the low priority system threads comprise an idler thread, wherein the idler thread is always ready to run.
- 20 17. Computer system comprising a processor, a memory accessible to the processor, an operating system and a computer program emulating a Java processor running on the processor, the computer program providing a scheduler for scheduling the running of Java threads on the processor, the scheduler maintaining a list of a plurality of threads, the plurality of threads comprising at least one high priority system thread and at least one user thread and, for each of the
- 25 the plurality of threads, maintaining a timestamp value, a service fraction value and a time left value; the scheduler adapted to:
- 30 a) identify as a current thread one of the plurality of threads which is ready to run and has a timestamp value earlier than the timestamp value of any other of the plurality of threads which is ready to run;
- 35 c) if the time left value for the current thread is greater than zero, i) transfer control to the current thread for a running time not exceeding the time left value for the current thread; and,

- 25 -

- ii) if the current thread is not a system thread, subtract the running time from the time left value for the current thread; and,
 - d) if the time left value for the current thread is not greater than zero, assign a new timestamp value to the current thread, the new timestamp value being later than the timestamp values of any of the user threads, and initializing the time left value of the current thread.
- 10 18. A computer readable medium having computer readable program logic recorded thereon, the program logic, when run on a computer, implementing a method comprising:
 - 15 a) maintaining a list of a plurality of threads, the plurality of threads comprising at least one high priority system thread and at least one user thread and, for each of the plurality of threads, maintaining a timestamp value, a service fraction value and a time left value;
 - b) identifying as a current thread one of the plurality of threads which is ready to run and has a timestamp value earlier than the timestamp value of any other of the plurality of threads which is ready to run;
 - 20 c) if the time left value for the current thread is greater than zero,
 - i) transferring control to the current thread for a running time not exceeding the time left value for the current thread; and,
 - 25 ii) if the current thread is not a system thread, subtracting the running time from the time left value for the current thread; and,
 - d) if the time left value for the current thread is not greater than zero, assigning a new timestamp value to the current thread, the new timestamp value being later than the timestamp values of any of the user threads, and initializing the time left value of the current thread.
- 30
- 35 19. A computer system comprising:
 - a) a processor;

- 26 -

- b) a multi-threaded operating system running on the processor, the operating system comprising a plurality of system threads;
- c) user software running on the computer system, the user software comprising one or more user threads;
- 5 d) a memory accessible to the processor, the memory containing a data structure comprising a record for each of a plurality of threads, the plurality of threads comprising at least one high priority system thread and the one or more user threads, each record comprising a timestamp value, a service fraction value and a time left value for a thread corresponding to the record;
- 10 e) a scheduler, the scheduler adapted to identify as a current thread one of the plurality of threads which is ready to run and has a timestamp value earlier than the timestamp value of any other of the plurality of threads which is ready to run and, if the
- 15 time left value for the current thread is greater than zero,
 - i) transfer control of the processor to the current thread for a running time not exceeding the time left value for the current thread; and,
 - ii) if the current thread is not a system thread, subtracting
 - 20 the running time from the time left value for the current thread; or,
- if the time left value for the current thread is not greater than zero, assigning a new timestamp value to the current thread, the new timestamp value being later than the timestamp
- 25 values of any of the user threads, and initializing the time left value of the current thread.

1/4

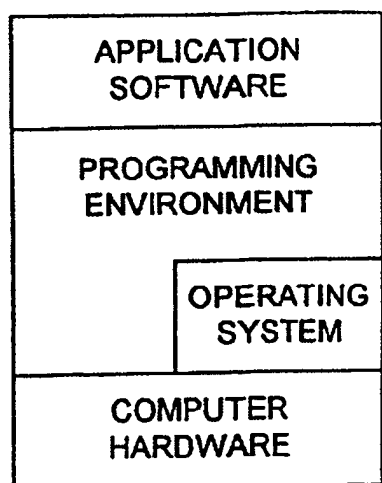


FIG. 1A
(PRIOR ART)




FIG. 1B
(PRIOR ART)








FIG. 1C
(PRIOR ART)

2/4

**FIG 2**

3/4

**FIG. 3****FIG. 5****FIG. 6**

4/4

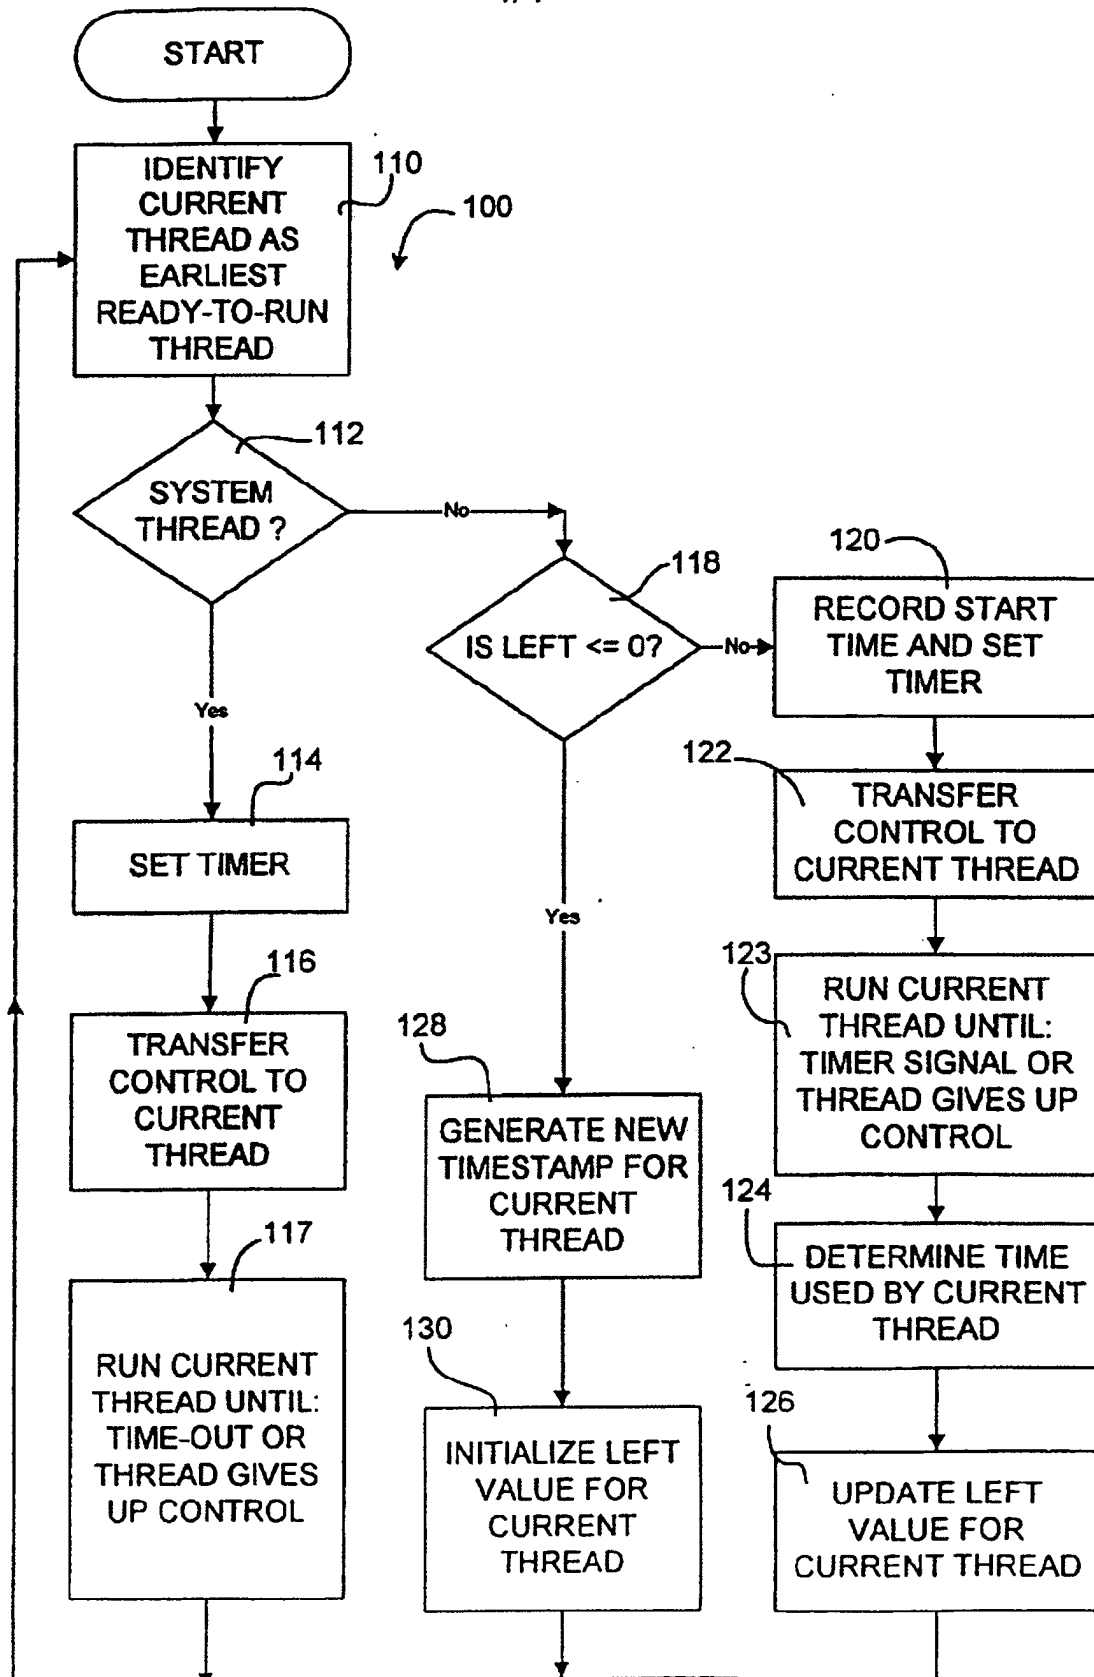


FIG. 4

*** Slip Sheet ***

(19) World Intellectual Property Organization
International Bureau(43) International Publication Date
13 June 2002 (13.06.2002)

PCT

(10) International Publication Number
WO 02/046980 A3(51) International Patent Classification⁷: G06F 17/60

(21) International Application Number: PCT/DK01/00801

(22) International Filing Date:
30 November 2001 (30.11.2001)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
60/251,862 8 December 2000 (08.12.2000) US(71) Applicant (for all designated States except US): CON-
FIGIT SOFTWARE APS [DK/DK]; Glentevej 67,
DK-2400 Copenhagen NV (DK).

(72) Inventors; and

(75) Inventors/Applicants (for US only): LICHTENBERG,

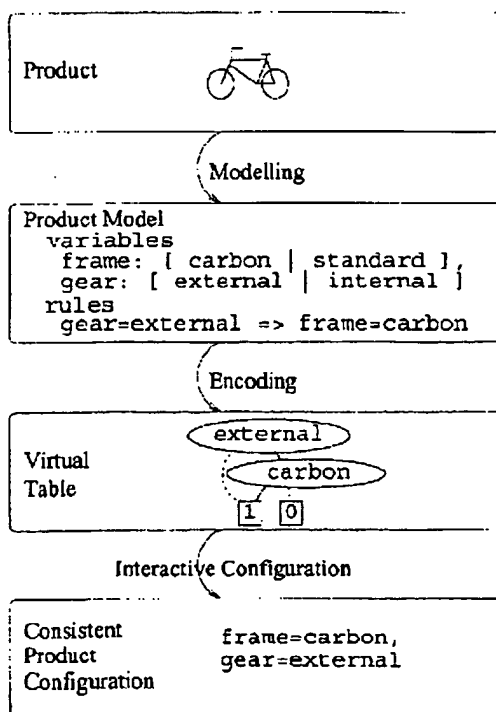
Jakob [DK/DK]; Lindegaardsvej 12 A, 1. tv., DK-2920
Charlottenlund (DK). ANDERSEN, Henrik, Rolf
[DK/DK]; Klirevænget 55, DK-2880 Bagsværd (DK).
HULGAARD, Henrik [DK/DK]; Koldinggade 20, 1.
tv., DK-2100 Copenhagen Ø (DK). MØLLER, Jes-
per [DK/DK]; Frederikssundsvej 14a 4. mf., DK-2400
Copenhagen NV (DK). RASMUSSEN, Anders, Steen
[DK/DK]; Nørgaardsvej 18b, 2. th., DK-2800 Lyngby
(DK).

(74) Agent: PLOUGMANN & VINGTOFT A/S; Sankt Anne
Plads 11, P.O. Box 3007, DK-1021 Copenhagen K (DK).

(81) Designated States (national): AE, AG, AL, AM, AT, AU,
(utility model), AU, AZ, BA, BB, BG, BR, BY, BZ, CA,
CH, CN, CO, CR, CU, CZ, CZ (utility model), DE, DE
(utility model), DK, DK (utility model), DM, DZ, EC, EE,
EE (utility model), ES, FI, FI (utility model), GB, GD, GE,
GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR,
KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK,
MN, MW, MX, MZ, NO, NZ, OM, PH, PL, PT, RO, RU,

[Continued on next page]

(54) Title: A METHOD OF CONFIGURING A PRODUCT USING A DIRECTED ACYCLIC GRAPH



(57) Abstract: A complex product is composed of several parts, where each part may depend on the other. As a consequence of these inter-dependencies, the selection of one part might exclude other parts from being included in the finished product. A *consistent configuration* is a selection of parts where all inter-dependencies are satisfied. A computer program for computer-assisted configuration helps an end-user to make choices that will lead to a consistent product. The preferred embodiment of the present invention, *Virtual Tabulation*, is a method for keeping track of inter-dependencies a large number of parts, to allow for the construction of an efficient *and* exact configuration program. Such a program allows interactive configuration over networks (e.g., the Internet). Another aspect of the invention, called *Smart Search*, allows a set of inter-dependencies among parts to be computed from a product database.

Published:

— with international search report

(88) Date of publication of the international search report:

26 September 2002

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

INTERNATIONAL SEARCH REPORT

International Application No.

PCT/DK 01/00801

A. CLASSIFICATION OF SUBJECT MATTER
IPC 7 G06F17/60

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 7 G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the International search (name of data base and, where practical, search terms used)

WPI Data, INSPEC, EPO-Internal

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	SABIN D ET AL: "PRODUCT CONFIGURATION FRAMEWORKS - A SURVEY" IEEE EXPERT, IEEE INC. NEW YORK, US, vol. 13, no. 4, 1 July 1998 (1998-07-01), pages 42-49, XP000791757 ISSN: 0885-9000 the whole document	1-42
Y	WO 00 13113 A (ANDERSSSEN H R ET AL) 9 March 2000 (2000-03-09) page 2, line 18 - line 25 abstract the whole document	1-42

	--- --/--	

☒ Further documents are listed in the continuation of box C.☒ Patent family members are listed in annex.

* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier document but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"S" document member of the same patent family

Date of the actual completion of the international search

12 June 2002

Date of mailing of the international search report

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3016

Authorized officer

Jenny Forss

INTERNATIONAL SEARCH REPORT

International Application No

PCT/DK 01/00801

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	ATTARDI G ET AL: "Web-based configuration assistants." AI EDAM ARTIF. INTELL. ENG. DES. ANAL. MANUF., vol. 12, no. 4, September 1998 (1998-09), pages 321-331, XP002902516 CAMBRIDGE UNIV. PRESS (UK) ISSN: irm-issn 0890-0604. the whole document ---	1-9,25, 26, 32-37, 39-42
A	CHAN W ET AL: "MODEL CHECKING LARGE SOFTWARE SPECIFICATIONS" IEEE TRANSACTIONS ON SOFTWARE ENGINEERING, IEEE INC. NEW YORK, US, vol. 24, no. 7, 1 July 1998 (1998-07-01), pages 498-519, XP000781149 ISSN: 0098-5589 the whole document ---	10-42
A	US 5 825 651 A (GUPTA N ET AL) 20 October 1998 (1998-10-20) column 2, line 50 - line 60 column 5, line 67 -column 6, line 6 abstract; figure 5 -----	1-9,25, 26, 32-37, 39-42

INTERNATIONAL SEARCH REPORT

Information on patent family members

International Application No

PCT/DK 01/00801

Patent document cited in search report		Publication date	Patent family member(s)	Publication date
WO 0013113	A	09-03-2000	AU 5369999 A	21-03-2000
			WO 0013113 A1	09-03-2000
			EP 1105821 A1	13-06-2001

US 5825651	A	20-10-1998	AU 731180 B2	22-03-2001
			AU 4236897 A	26-03-1998
			BR 9712795 A	14-12-1999
			EP 0925543 A1	30-06-1999
			JP 2000517450 T	26-12-2000
			NO 990864 A	15-04-1999
			WO 9810360 A1	12-03-1998
			US 6405308 B1	11-06-2002

*** Slip Sheet ***

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
13 June 2002 (13.06.2002)

PCT

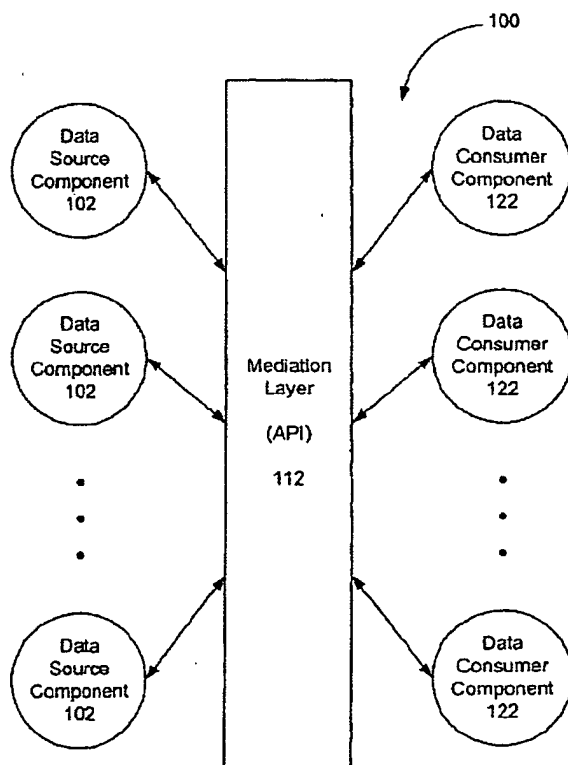
(10) International Publication Number
WO 02/46916 A2

(51) International Patent Classification⁷: G06F 9/00
(21) International Application Number: PCT/US01/48418
(22) International Filing Date: 22 October 2001 (22.10.2001)
(25) Filing Language: English
(26) Publication Language: English
(30) Priority Data:
60/242,041 20 October 2000 (20.10.2000) US
(63) Related by continuation (CON) or continuation-in-part (CIP) to earlier application:
US 60/242,041 (CIP)
Filed on 20 October 2001 (20.10.2001)
(71) Applicant (for all designated States except US):
POLEXIS, INC. [US/US]; 2815 Camino del Rio South,
San Diego, CA 92110 (US).

(72) Inventors; and
(75) Inventors/Applicants (for US only): KADEL, Richard, William, Jr. [US/US]; 14524 Rutledge Square, San Diego, CA 92128 (US). HERMAN, Jeffrey, Stephan [US/US]; 1843 Narragansett Court, San Diego, CA 92107 (US). EXLINE, Christopher, Lee [US/US]; 8215 Royal Gorge Drive, San Diego, CA 92119 (US). ALMILLI, David, Edward [US/US]; 4818 Cypress Street, #4, La Mesa, CA 91941 (US). PRIEBE, Christopher, C. [US/US]; 4525 Gesner Street, San Diego, CA 92117 (US).
(74) Agents: HALL, David, A. et al.; Heller Ehrman White & McAuliffe LLP, 4350 La Jolla Village Drive, 6th Floor, San Diego, CA 92122-1246 (US).
(81) Designated States (national): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC.

[Continued on next page]

(54) Title: EXTENSIBLE INFORMATION SYSTEM (XIS)



(57) Abstract: A framework enables data source components to be developed independently of data consumer components. A mediation layer, typically implemented as a group of APIs (application programming interface), handles and defines the mediation and interface between the source and data components. The framework, called XIS (extensible information system), is especially suited for development of information-handling systems and applications. Data source components and data consumer components are typically designed to communicate with each other via several interfaces. Domain, relationship, attribute/metadata, and change event interfaces are defined within the mediation layer. Other interfaces may also be defined. Data source components that are written for non-XIS aware environments or frameworks may still be used with XIS by "wrapping" such source components with code to conform to the interface requirements. Java objects are examples of data source components. Data consumer components thus are able to use or consume various source components regardless of the data types and the data source. Thus, once a data consumer component is developed within the XIS framework, any data source components within the XIS framework may be consumed by a data consumer component.



WO 02/46916 A2



LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, PH, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, US, UZ, VN, YU, ZA, ZW.

- (84) **Designated States (regional):** ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Declaration under Rule 4.17:

- as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii)) for the following designations AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC,

EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, JS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, PH, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, UZ, VN, YU, ZA, ZW, ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG)

Published:

- without international search report and to be republished upon receipt of that report

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

EXTENSIBLE INFORMATION SYSTEM

A portion of the disclosure of this patent document contains material, which is subject to copyright protection. The copyright owner has no objection to the facsimile reproduction by anyone of the patent document or the patent disclosure, as it appears in the Patent and Trademark Office patent file or records, but otherwise reserves all copyright rights whatsoever.

This application claims priority of co-pending U.S. Provisional Patent Application Serial No. 60/242,041 entitled "Extensible Information System (XIS)" by R. Kadel et al., filed October 20, 2000. Priority of the filing date of October 20, 2000 is hereby claimed, and the disclosure of said Provisional Patent Application is hereby incorporated by reference.

BACKGROUND

1. Field Of The Invention

The present invention relates generally to data processing systems and, more particularly, to systems that process, analyze, and display data or information.

2. Description of the Related Art

The continuing inroads made by computer technology into the practices of information storage and manipulation have opened up new realms of possibility for intelligent, informed decision-making. From labor statistics and scientific databases like the human genome project to traffic patterns and aircraft positions, the

availability of information in electronic form allows sophisticated analysis techniques and display methods to be applied at the touch of a button. However this diverse array of information also poses significant challenges in areas such as:

- effective organization of information;
- 5 - ability to interpret and manipulate information in an increasingly wide variety of formats so that data and processing software can be brought together in a compatible but also timely and effective way; and
- recognizing and navigating relationships between disparate
- 10 information types.

This task of information management has already grown beyond the ability of human operators to keep pace with it, and in most applications, there is far more relevant information electronically available on the web and elsewhere than is effectively used. The problem is that software affording display of and interaction

15 with information must be custom-designed for the particular type and format of the information it will work with. This leads to high costs of software development and limited availability of software appropriate to information of interest, and poor integration between information from different sources. An analyst must use one program when dealing with geographic distribution data such as precipitation amounts

20 or land use, another for rendering 3-dimensional terrain, another for examining congressional districts and voting records, and so on. Many desirable applications of available data, such as determining correlations between political party and land use strategy in this case, fall between the cracks left by specialized applications, and there is no easy way to allow them to interoperate.

25 The primary means for allowing one application to operate on the results of another are the limited cut and paste facilities offered by many operating systems, which allow plain or formatted text or images to be moved between programs. There

is no independent way to transfer structured data such as a table of numbers with particular column headings, a set of lists of items under specified categories, or even something as simple as a single number with an associated unit of measurement.

From the discussion above, it should be apparent that there is a need for
5 computer systems that have greater ability to develop, integrate, and interoperate with disparate sources of information, to more easily develop software applications, components, or objects, and to facilitate interoperation of data between software components. The present invention fulfills this need.

10 SUMMARY OF THE INVENTION

In accordance with the present invention, a framework for use with object oriented programming (OOP) systems provides a framework user with classes that comprise a mediation layer that defines an interface between data source components
15 and data consumer components such that the configuration of the data source components can be specified independently of the data consumer components. The mediation layer specifies data relationships of the data objects, domain methods for defining groups of class attributes, attribute metadata for defining groups of class attributes, and change event registration for detecting changes in data values. Thus,
20 when extended, the mediation layer of the framework can support runtime data manipulation between unrelated data source components and data consumer components. In this way, the framework provides an extensible information system. The framework will be referred to herein as "XIS", and is especially suited to assist in the development of information-handling systems or applications.

25 Data source components that are configured for a non-XIS-aware programming environment or framework may still be used with XIS by "wrapping" such source components with code to conform to the interface requirements. Data

objects of the well-known "Java" programming language are examples of data source components. Data consumer components thus are able to use or consume various data source components regardless of the data types and the data source. Thus, once a data consumer component is developed within the XIS framework, any data source
5 components within the XIS framework may also accordingly be used. The framework can be provided as a group of APIs (application programming interface).

The XIS framework of the present invention also includes libraries and APIs that provide information management technologies that enable developers and integrators to combine XIS-aware components, data sources, and off-the-shelf
10 JavaBeans into complete systems designed around whatever architecture is best for the situation. XIS further supports multiple, dynamic domains and is scalable to n-tier systems using the latest technology. The framework enables many developers to choose architectures based on their requirements (e.g., server, client/server, application server, web server, standalone, hand-held, etc.).

15 Other features and advantages of the present invention should be apparent from the following description, which illustrates, by way of example, the principles of the invention.

BRIEF DESCRIPTIONS OF THE FIGURES

20

Figure 1 is a diagram of the extensible information system (XIS) framework constructed in accordance with the present invention.

Figure 2 is a more detailed diagram of Figure 1, showing the various interfaces within XIS, constructed in accordance with the present invention.

25

Figure 3A is a more detailed diagram of Figure 1, including information and services available within XIS, constructed in accordance with the present invention.

Figure 3B is a hierarchical structure diagram of an embodiment of an InfoModel constructed in accordance with the present invention.

Figure 4 is a more detailed diagram of Figure 1 showing various information available to components within XIS constructed in accordance with the present invention.

Figure 5 is a diagram of the semantic representation of a data item within the XIS framework constructed in accordance with the present invention.

Figures 6A to 6F show data consumer components, particularly display components, constructed in accordance with the present invention.

Figure 7 shows a consumer component displaying the properties exposed by a data source object constructed in accordance with the present invention.

Figure 8 is a unified modeling language (UML) diagram of a Person object.

Figure 9 lists a set of attributes exposed within the XIS framework by the Person object of Figure 6 constructed in accordance with the present invention.

Figure 10 illustrates that an object constructed in accordance with the present invention may subscribe to more than one domain policy or definition.

Figure 11 shows a domain usage and sequence scenario between a consumer component and source component constructed in accordance with the present invention.

Figures 12A to 12D list information regarding a Type Metadata package class typically implemented in the mediation layer and constructed in accordance with the present invention.

Figure 13 shows a relationship usage scenario diagram between a consumer component and source component constructed in accordance with the present invention.

Figure 14 is a flow chart that shows a reference resolution scenario usage sequence in which an indirect reference to an information element is created, passed

around, and then resolved by two separate data consumers, causing reconstruction to occur only the first time.

Figure 15 lists an exemplary Java source file to implement an information-handling application constructed in accordance with the present invention.

5 Figures 16A to 16B list an exemplary Java source file to implement an information-handling application constructed in accordance with the present invention.

Figure 17 shows a data consumer component using a source component constructed in accordance with the present invention.

10 Figures 18A to 18B list an exemplary Java source file to implement an information-handling application constructed in accordance with the present invention.

15 Figures 19A to 19B list an exemplary Java source file to implement an information-handling application constructed in accordance with the present invention.

Figure 20 shows a data consumer component using a source component constructed in accordance with the present invention.

Figure 21 lists an exemplary Java source file to implement an information-handling application constructed in accordance with the present invention.

20 Figures 22A to 22C list an exemplary Java source file to implement an information-handling application constructed in accordance with the present invention.

25 Figures 23A to 23D list an exemplary Java source file to implement an information-handling application constructed in accordance with the present invention.

Figures 24A and 24B show two data consumer components using a source component constructed in accordance with the present invention.

Figures 25A to 25C list an exemplary Java source file to implement an information-handling application constructed in accordance with the present invention.

Figure 26 shows a data consumer component using a source component
5 constructed in accordance with the present invention.

Figure 27A and Figure 27B show a flow chart of a data exposure facility within the XIS framework constructed in accordance with the present invention.

Figure 28 is a diagram of how InfoModels constructed in accordance with the present invention provide contextualization with the XIS framework.

10 Figure 29 is a diagram of the pluggable service facilities within XIS constructed in accordance with the present invention.

Figure 30 shows a flow diagram of a wizard or an API constructed in accordance with the present invention that assists in creating XML (extensible markup language) DSIs and database DSIs.

15 Figure 31 shows objects and methods involved in distribution collaboration facilities constructed in accordance with the present invention.

Figure 32 shows semantic repositories for distributed agents constructed in accordance with the present invention.

Figures 33A to 33P list information regarding a ContentInfoBean class
20 constructed in accordance with the present invention.

Figures 34A to 34D list an information management package typically implemented in the mediation layer and constructed in accordance with the present invention.

Figures 35A to 35F list an InfoModel interface package typically implemented
25 in the mediation layer and constructed in accordance with the present invention.

Figures 36A to 36B list a package for handling change events typically implemented in the mediation layer and constructed in accordance with the present invention.

Figures 36C and 36D show two sequence diagrams showing how change events are handled and constructed in accordance with the present invention.

Figures 37A to 37D list an exemplary Java source file to implement an information-handling application constructed in accordance with the present invention.

Figures 38A, 38B, 38C, and 38D show two data consumer components using a source component constructed in accordance with the present invention.

Figures 39A, 39B, and 39C is list information concerning an AttributeAlias class of the framework.

Figure 40 is a block diagram of a computer device that may be used to operate with the framework, in accordance with the present invention.

15

DETAILED DESCRIPTION

The following detailed description illustrates the invention by way of example, not by way of limitation of the principles of the invention. This description will clearly enable one skilled in the art to make and use the invention, and describes several embodiments, adaptations, variations, alternatives and uses of the invention, including what we presently believe is the best mode of carrying out the invention.

The invention will be described by way of illustration with reference to various classes, objects, sample codes, etc. written within the exemplary framework, but it should be understood that such classes, class libraries, application programming interfaces, interfaces, objects, etc. may be differently coded, implemented, designed, etc. and yet support the functions and features of the present invention.

25

Object-Oriented Technology

Many application programs and APIs (application program interface) are developed using object-oriented (OO) technology. Using OO technology, a system is typically developed as a collection of interrelated cooperative objects that are
5 instances of classes that typically include corresponding states and behaviors. A class is a blueprint or template that defines the variables and the methods common to all objects that are instances of the class. An object maintains its state in one or more variables and implements its behavior with methods or functions.

Object-oriented programming (OOP) techniques encapsulate, or bind together,
10 data and the methods that operate on them. This encapsulation permits program development to more closely model real-world entities and breaks up program development efforts into smaller, more manageable pieces. Although OOP techniques have done much to improve program development efficiency, such techniques still require a great degree of code generation on the part of developers,
15 which discourages program reuse.

Even with OOP techniques, specifying data sources and the way in which data will be retrieved or updated from the data sources are still typically hard-coded within the objects themselves. An object, for example, may be specified to display pricing data (for example, a price display object) and would likely be hard-coded such that
20 the data source is predefined. Thus, an object may be written to retrieve data from a source such as a relational database management system (RDBMS), a spreadsheet file such as a spreadsheet in the format of "EXCEL" spreadsheet application by Microsoft Corporation, or the object may be written to an XML file. Nevertheless, in conventional systems, the price display object would be limited such that if a different
25 source of data, for example, an XML file rather than an EXCEL file, has to be used, the object has to be modified to incorporate access to and manipulation of the XML file. Furthermore, if a different set of fields are retrieved, for example, latitude and

longitude (i.e., data fields unrelated to pricing information), a new object to display latitude and longitude may have to be written to display such information in tabular format.

Table I below shows the data configuration of a typical object.

5

Table I
State/Properties: Name, SSN, position, DOB, and . . .
Behavior: Record[]
getRecords(DateRange)
promote ()
. . .

For an application to be able to use this data object, the application would need prior knowledge of object properties (state) and behavior. The application would be tightly coupled with the class from which the object was instantiated. The meaning and intent of the data might be found in the source code or by asking the original programmers (e.g., SSN is social security number stored as string rather than a numeric type).

Assuming the object was created by one application written independently of another, the object could not be intelligently exchanged between the first application and the other with the expectation that the object would be processed unless the applications were tightly coupled. A way to have objects be used in various applications is thus highly desirable to reduce programming time and resources.

Common OOP languages include "Java" from Sun Microsystems, Inc. of Palo Alto, California, USA, C++, Simula, and Smalltalk.

Framework

The concept of a framework is an important part of OO programming technique. A framework is a specification of the classes and the relationships between classes such that the framework defines a class hierarchy that can be used
5 over and over again, with overrides and extensions. In this way, an initial problem solution specified by a class hierarchy can be adapted and customized for new circumstances, simplifying program maintenance. In essence, a framework is a set of OOP classes that embodies a predetermined set of attributes and methods for providing a common group of behaviors.

10 OOP frameworks have been developed in an effort to further reduce program development costs. An application program developer utilizes the framework and builds upon it, starting with the classes, attributes, and methods defined by the framework designer and adding subclasses and attributes and modifying methods depending on the problem to be solved. Such changes to the framework are typically
15 referred to as framework extensions, and are made possible by the OOP notions of inheritance and polymorphism. Thus, a framework can speed the development of an OO application program. The challenge confronting framework developers, then, is to define a set of classes and methods that best supports the desired problem solution and will accept the most likely framework extensions. Thus, the designer of a
20 framework must carefully assess what framework users will most likely need in the way of classes, attributes, and methods.

It is therefore a technical advantage of the present invention to provide a framework that allows information-handling software to adapt to new data types and formats of information, so that a single application built within this framework has a
25 significantly broader range of compatibilities and domains of usefulness than a conventional application and, further, may automatically interoperate at a structured level with other similarly enhanced applications. In one embodiment of the present

invention, desirable applications of data such as determining correlations between political party and land use strategy may be facilitated.

In an embodiment of the present invention, a single information visualization or manipulation application is able to handle many diverse types of information, even
5 those that are conceived and developed subsequently to the completion of development of the application itself.

In another embodiment of the present invention, information from multiple and diverse sources yet sharing some basic feature in common - such as being distributed in geographic space, or being distances measured in meters - can be simultaneously
10 displayed (superimposed) and manipulated within a single application. This invention further provides means of access to common data sources such as relational databases, extensible markup language (XML) streams, and "Java" programming language software objects.

In an embodiment of the present invention, an information-handling
15 application is defined as any software application which is capable of utilizing digitally available structured information in bulk form, such as from electronic files, databases, or internet web sites to provide a display and/or a set of possible manipulations to a user, such as after some internal processing and transformation of the information such as summarizing it, performing computations on it, or selecting a
20 subset. Examples include but are not limited to graphing programs, analysis tools (for handling financial data, statistical, time series, etc.), and interfaces to geographic information systems. Prototypical information-handling applications process structured information that has some form of hierarchical and/or modular structure. In one embodiment, image and word processing applications are not examples of
25 what we term information-handling applications. In another embodiment of the present invention, structured information in bulk form does not include plain text or specific electronic media files such as MP3 files.

A conventional information handling application such as a graphing program comprises two major components: one component handles the intake of information, from storage devices, the network, or user input, and the other component handles the display of the information. There may also be parts of the application that perform
5 computations or transformations on the information before it is displayed. From the developer's perspective, a modular piece of software that handles a portion of (or all of) the data intake function is a data source component, because it provides information to other components within the application. A modular piece of software that handles part of (or all of) either the computation, transformation, or display of
10 information is a data consumer component, because it uses (consumes) information provided by the data sources.

Data sources and data consumers communicate with each other by making function calls. The particular functions and their parameters are defined as part of an internal version of an application program interface (API). The data consumer
15 components handle data that is given them through specific calls on their respective API, so that information provided through data source components written for different applications will not be accepted. This situation provides little opportunity for reuse of data source or consumer components, since each is written to interoperate only with a specific instantiation of the other.

In an embodiment of the present invention, an internal mediation layer is inserted between the data source components and the consumer components that expose data structure in a standardized way. Consumer components are constructed to use this common API, which is suitable for expressing an extremely wide variety of information types. Data source components can be constructed to deliver the
20 information they take in this format, or a small amount of code can be written to translate the output of another source component into the common format. The result of the construction is that any consumer component can work with any data source

component regardless of the specific nature of the information involved, or whether one was anticipated during the design of the other. The common API allows consumer components to automatically extract whatever features of information from a particular source they are most suitable for displaying or computing with, and also
5 allows them to utilize all the other features of that information in a generic, unspecialized fashion.

Other embodiments of the present invention comprise: (1) a system for representing information such that: (a) a single fixed interface suffices to describe a wide range of information types, (b) the information is rendered self describing to an
10 extent, and (c) relationships between different information elements are expressed; (2) an apparatus for allowing this information representation to be employed as a medium between data source and data consumer software components; (3) a method and apparatus for attaching clarifying material on "intended use" to information so that consumer components can handle it more appropriately; (4) a method and apparatus
15 for providing a context to all information consumption affording control over security and visibility of data; (5) a method and apparatus for allowing users to transfer information from one consumer component to another through intuitive "drag and drop" and "cut and paste" interfaces; (6) a method and apparatus for developing enhanced information handling applications based on the foregoing framework; (7) a
20 method and apparatus for automatically re-representing data from partially self describing sources including relational databases, XML streams, and Java software objects; (8) a method and apparatus allowing end users using computers distributed over a network to collaboratively view and manipulate information; and (9) a method and apparatus allowing software components distributed over a network to
25 automatically obtain annotations on intent and other aspects of encountered data.

Most aspects of this invention can be implemented in any object oriented programming language such as Smalltalk, Objective C, C++, or Java. Certain aspects

of it, pointed out below, are especially suited to object oriented languages such as Java and Objective C that provide for run time self analysis by programs. However, these are in every case peripheral aspects, and the essential parts of the invention may be implemented in an object oriented language like C++ without these capabilities.

- 5 We will from time to time make reference to a "preferred embodiment" implemented in Java, but this should not be taken to be limiting.

In a preferred embodiment the present invention provides a system and/or framework of software components for purposes of aiding development of information-handling applications. The system and/or framework of object-oriented software components aid the development of applications that read, gather, or receive
10 electronic information and allow the display of such information to and manipulation by users. These applications are structured in a way so that there are distinct data source components and data consumer components. The crucial features of the system and/or framework are the fact that data is seen by consumers only through a standard, conventionalized interface that incorporates: a breakdown into attributes,
15 relationships, semantically-annotated "domain methods", event-broadcasting of changes in the foregoing to all registered consumers, metadata for each attribute providing certain declarative and procedural information including unit of measure, content length, data quality, default value, comparator function, summary function, validation function, and input/output functions for display and user editing within an
20 extensible set of interface modalities, related data items available for each data item including those in child or generic relationships, and including the ability to store a reference without direct access to the related item itself (it can be reconstructed if needed), and contextualization such that the attributes, relationships, events, and
25 methods available for a data item depends on context (see below).

In a related embodiment, the system and/or framework includes methods for data exposure wherein the system determines automatically which method is

employed in any particular case, and the choice is invisible to consumers. Exemplary methods such as, a data source component which directly provides the standard interface, a data source component which is accompanied by a separate Translator component which maps its interface to the standard interface; and a data source
5 component which is automatically inspected by the system to determine what available data fields it has, these are exposed to consumers through the standard interface as detailed above.

In a preferred embodiment, the contextualization provided by the system and/or framework is accomplished by delegating the final responsibility for data
10 exposure to an object termed an "InfoModel" which is able to choose which of the available attributes and methods on an object should be exposed or hidden, and is able to add additional attributes or methods. Preferably, the system and/or framework further provides facilities supporting transfer of data items between consumers, either by user manipulation (cut/paste or drag/drop) or by internal method calls. When a
15 transfer is initiated, the data item is passed without its contextual characteristics (supplied by an InfoModel) to the receiver. The system and/or framework includes facilities supporting the automatic integration of service components (generally comprising management facilities such as menu provision or running of user dialog routines and may or may not have data source provider and/or user interface provider
20 components) into an application. Integration is handled by a plug-in manager object and set of interfaces such that: each service provider is given the opportunity upon loading to query all components (data providers or consumers) existing within the application for whether they desire the service, and if so what parameters they would like to pass to it. The provider then acts on these responses. Whenever new client
25 components are loaded, they will be queried and possibly responded to by all services loaded within the application.

In another embodiment a system for importing partially schematized or structured data sources (such as a database system or an XML document) into a system as described herein. Any application written using the disclosed system and/or framework can then utilize these data sources through a standard interface. In particular, a system and/or framework for performing all or a substantial portion of the functionality outlined herein - if it is in connection with a framework or application implemented using a framework falling under categories as disclosed such as an infrastructure for mediating between information sources and consumers and rapidly assembling applications using them.

Thus the present invention provides a method and apparatus involving two subsystems, one which processes the schema information for a data source to determine the type and metadata information for information attributes from the source, and one which processes the instance information. Preferably, the schema subsystem processes schematic information such as the type information contained in an XML schema document or the column information available from a relational database management system and sets up appropriate metadata and relationship structure for information elements, but it does not necessarily create any information elements. Optionally, additional annotations can be provided by a user in the form of a file specifying how to map attributes in the original data to metadata, relationships, or domain policy attribute or method definitions. Preferably, the instance subsystem processes instance information such as that available from an XML document or set of relational database rows, for which corresponding schema information can be found, and constructs information elements exposed through the standard interface. Data consumers can utilize the outputs of these as if they were custom-created data sources for the type of information described in the schemas.

In a related aspect, implementation of the invention as described herein for both XML and relational databases is provided. In the XML implementation, a built-

in data type hierarchy for XML schemas is used to provide types to attributes, and the element-subelement-attribute structural hierarchy in XML schemas is used to determine relationships between information elements. All documents and schema references to other documents are followed up to their sources in order to specify further attribute or relationship information. A user can optionally specify an XSLT (XML Stylesheet Transformations) document that maps the typed data fields found from the instance-schema combination into attributes with richer metadata and/or references to specific domain policies. In the database implementation, the table column heading types from the relational database (date, integer, text, etc.) are used to provide default type metadata to the attributes of information elements derived from table rows of the database. A user can optionally specify a mapping between columns or sets of columns into attributes with specific metadata. Furthermore, a user can define a sequence of queries for which the results are to result in a hierarchy of information elements. For example, for each information element retrieved from, for example, Query 1, a parameterized instantiation of Query 2 may be applied to retrieve elements which will be exposed as the children of the first information element. For each of these elements, a parameterized instantiation of Query 3 may be instantiated, and so forth.

Framework Block Diagram

Figure 1 shows a basic block diagram of a framework that is implemented in accordance with the present invention. The framework 100 of Figure 1 can be used to develop, for example, an information-handling system or application wherein data source components 102 and data consumer components 122 are separate and independent from each other but can communicate and share data. The framework 100 described herein shall be called the extensible information system (XIS) framework.

From a developer's perspective, a modular piece of software that handles part or all of the data intake or retrieval function is a data source component 102. A modular piece of software that handles part or all of either data computation, transformation, presentation, or display of information is a data consumer component 122. A conventional information-handling application such as a graphing program comprises two major components: one side handles the intake of information from storage devices, the network, or user input, and the other side handles the display or presentation of the information. There may also be parts of the application that perform computations or transformations on the information before it is displayed.

10 XIS Framework

In the XIS framework 100 constructed in accordance with the invention, there may be more than one data source component 102 (also referred to as a source, data source, source component, and source object) and there may be more than one data consumer component 122 (also referred to as a consumer, consumer object, consumer component, and data consumer). The data source components 102 provide information to other components within the XIS framework 100 via a mediation layer 112, while the data consumer components 122 use the provided information.

The source components 102 and consumer components 122 that comprise the XIS framework 100 are configured such that they communicate with each other in a manner that is supported by the programming environment in which the framework exists, typically by making function calls, via the mediation layer 112. In one embodiment, the mediation layer 112 consists of a group of application programming interfaces (APIs) or class libraries that define a common format for data exchange, establishing an information exchange standard. Thus, data source components 102 configured for a different system, application, or framework may still be used within the XIS framework if they are "wrapped" in a predetermined way or are suitably

modified, as will be known to those skilled in the art, such that they conform to the requirements of the data exchange interface of the mediation layer 112.

A relatively simple block of programming code may be produced to translate the output of a data source component 102 into the common format as defined by the mediation layer or API 112. In this way, any consumer component 122 may work with any data source component 102 regardless of the specific nature or type of information involved and regardless of whether the data source component was known or anticipated during the design of the data consumer component. Because of the XIS framework described herein, the data consumer 122 may also extract whatever features or sets of information that it needs or can process from one or more data source, thereby enabling the data consumer to utilize features and information in a generic unspecialized fashion.

The mediation layer 112, which is between the data source components 102 and the consumer components 122, exposes the data structure of the data source components 102 in a manner that is common for the framework, preferably through an API interface. Data consumer components 122 are constructed to use this common API, including utilizing the exposed data structure. The data consumer components 122 are suitable for expressing an extremely wide variety of information types. In addition, the data consumers may be constructed to deliver the information received from the mediation layer to another data consumer component, or to output the information received for display.

Data consumer objects 122 are application or software objects or components that use data, or may be said to consume it. Examples of data consumers include display programs that display data in tabular format, in graphical mode, in organizational chart mode, spreadsheet mode, timeline mode (similar to files in the format of the "PROJECT" application from Microsoft Corporation), and the like. Data consumers may be desktop based, web-based (Internet-based), distributed

architecture, interpreted programs, and the like. Data consumer components within the XIS framework may provide display, computational, or interactive facilities. In the present description, such specially configured data consumer components are referred to as "INFOBEAN" objects, as available from the assignee of the present
5 invention, Pollexis, Inc. of San Diego, California, USA. Those skilled in the art will understand how to extend the framework 100 to produce desired applications, in view of the description herein.

Figure 2 shows additional details of the mediation layer 112 in the framework 100. In the framework 100, four special types of interfaces may be exposed to the
10 data consumer components 122. These four interfaces include Domain Definition or Domain Policy 202, Relationship 210, Attributes/Metadata 214, and Change Event 216. In one embodiment, change events are registered within the XIS framework 100. These interfaces may be exposed, for example, through class and object libraries that may be developed for the XIS framework. With these pre-defined class and object
15 libraries and APIs, developers of data source components 102 and data consumer components 122 may independently develop their own source and consumer components, thus facilitating development of information-handling applications or systems. An example of a class that may be implemented or written in the mediation layer to facilitate development of source and consumer components is illustrated in
20 Figures 33A to 33P, which show details of a sample class called ContentInfoBean.

Domain definitions or policies 202 (also referred to as domain objects or simply domains) are optional within the XIS framework 100. If defined, the corresponding Domain methods 212 for such policies are exposed to the data consumers 122. Other information about the domain policy may also be exposed,
25 such as Attributes/metadata 214 and Relationships 210 for that particular Domain Policy. Attributes/Metadata 214, Change Events 216, and Relationships 210 may also be exposed within the XIS framework 100, even if domain policies are not defined.

Relationships show the relationship between data source components. Containment, hierarchies, and ad-hoc relationships, for example, may be shown. They may be expressed as "members of" and references. Relationships are thus recognized and may be obtained accordingly, as shown in Figure 14 described further
5 below.

Change events may be used to ensure that all consumer or source components that have an interest in another source component's data are notified when changes occur. Changes may be notified when attributes, members (or containers), or references (or referrers) are added, updated, or removed and when values are changed.

10 Once a data consumer component 122 obtains a reference to a data source 102, the data consumer component may obtain access to these interfaces 202, 210, 212, 214, 216 through method calls in accordance with the programming environment of the host computer system. The data consumer 122 obtains such a reference typically in preparation to consume or use information from at least a particular data source
15 102. Examples of such consumer component method calls include:

- get attributes;
- get member and container information elements;
- get referred and referring information elements;
- get unique ID and selection state;
- 20 - add and remove instance specific attributes;
- register to be informed of changes to attributes or relationships;
- and
- invoke method calls of a semantically annotated type, described
below.

25 These types of method calls will be familiar to those skilled in the art and familiar with the programming environment of the host computer system.

Figure 3A is a schematic diagram showing the information about data source objects that is exposed by the mediation layer 112 in Figures 1 and 2, as well as the services available within the XIS framework 100. In accordance with the invention, data source components are wrapped with an object that exposes the interface
5 information for the data source to data consumers through the mediation layer.

In the preferred embodiment, the wrapper around each data source is an object provided by a class called InfoModel. As shown in Figure 3A, a data source comprising a raw data item 351, via the APIs 112 as described above, offers a variety of information, as shown in the blocks 354, 356, 358, 360, 362, through which data
10 consumers may gain access to the data source. The information represented by the blocks 354, 356, 358, 360, 362 is exposed or may be contained within an InfoModel 350, which provides an information space for all information available in the XIS framework 100 for the wrapped data source component. In one embodiment of the invention, the data needed by data consumers are accessed via such InfoModels (such
15 as further illustrated in Figure 28). The wrapper provided by the InfoModel is called a LeifDataItem object 352. Thus, the InfoModel is an interface that permits data sources and data consumers of the XIS framework to independently share their data. In this way, the methods 360, domains 358, metadata 354, relationships 356, and other information of a data source object are thus exposed using the InfoModel object,
20 which may be said to exist in the mediation layer 112, which sets up contexts for data consumption. The XIS framework also provides additional data services through the mediation layer, such as the Plug-In Service 264 and BeanContext Service 266 indicated in Figure 3A. These services provide the functionality listed in their respective boxes 364, 366. Other data services may be provided, as desired.

25 The InfoModel object provides a context for information access, i.e., context is instantiated through the InfoModel concept. In this way, security may be enforced, considering that the conceptual or real "end-user" accessing the information is known

and thus the underlying data source may be forwarded correctly. Furthermore, through this InfoModel concept, an embodiment of the framework may be implemented such that attributes and methods may be added or removed by the framework user. Original attributes and methods in one context may also be
5 overridden.

An attribute is any property of an object. In one embodiment of the invention, an attribute may be coded with an "AttributeDescriptor". An AttributeDescriptor is provided by the object or its Translator to describe a property of the object using TypeMetaData, e.g., how to display it, how to edit it, its permitted ranges, and the
10 like. Translators map data source-specific structures and types to an XIS-generic representation. Translators are further discussed below. Attributes may also be defined using Java introspection or reflection, as discussed further below.

Dynamic attributes may also be exposed within the XIS framework 100. These attributes are attributes that are exposed during run time, such as, for example,
15 a database query wherein each column is an attribute revealed at run time. The mediation layer has a mechanism that captures such information. In the preferred embodiment, the XIS framework operates in conjunction with a database management application, so that the mediation layer utilizes run time database query mechanisms that generate metadata of the data object, such as data table column headings, data
20 types, and the like. Such mechanisms may be provided, for example, by SQL queries and the like. In one embodiment, data consumer applications using the XIS APIs are unaware of whether an attribute for a data item is obtained dynamically or statically, so that the process is transparent to the application user.

A domain is a collection of semantically meaningful attributes and method
25 signatures grouped by their common domain of interest. Examples of domains include a Display Domain (a domain that describes how to display information) or a

Geo Domain (a domain that specifies how to handle geographic information).
Domains are further discussed below.

The contextualization provided by the XIS framework system is accomplished by delegating the final responsibility for data exposure to an object termed an
5 "InfoModel", which is able to choose which of the available attributes and methods of an object should be exposed or hidden, and is able to add additional attributes or methods. Preferably, the system and/or framework further provides facilities supporting the transfer of data items between data consumers, either by user manipulation (cut/paste or drag/drop) or by internal method calls. This supports data
10 sharing among data consumers. When a transfer is initiated, the data item is passed without its contextual characteristics (which are supplied by an InfoModel object) to the receiver, i.e., only non-contextual information is transferred. In this way, the source data may remain unchanged.

As illustrated in Figure 3A, a FieldMetaData file 354, a type of metadata, is
15 used to represent sorting, subset, and visibility criteria for an attribute. This is further explained below in conjunction with Examples 1 to 5, wherein subsets of data are shown depending on whether "All Attributes" or "Preferred Attributes" are shown. In one embodiment of the invention, default behavior method of domain policy attributes may also be overridden. Within this framework, context-specific special attributes for
20 flagging whether a data item should be visible or not (e.g., can be used for display filtering, including the XIS value slider, tree checkboxes in the map view, etc.) may be defined. This may be achieved through the mediation layer by utilizing a user interface to modify display processing for the data item. Furthermore, a selection attribute that is a context-specific special attribute for flagging whether an object is
25 "selected" (and typically should be highlighted) may also be defined. Selected pieces of information are often shared among views (context-specific, though), and certain operations may be performed on the set of selected objects in a given context

The InfoModel objects as described above know how to return a data item (a LeifDataItem) given a corresponding raw data source identifier. The InfoModel also manages these data items. As indicated in Figure 3A, an InfoModel may contain one or more LeifDataItem objects, each of which provides an interface to a raw data item.

5 An application developer may use an InfoModel object 350 to get a LeifDataItem object 352 in order to use the XIS APIs for accessing the underlying data source in a generic fashion. If two different data consumer components ask for a data item (e.g., a LeifDataItem) from the same InfoModel for the same raw data source, then both data consumers will get the same instance of the LeifDataItem object. InfoModel
10 objects may also be nested as indicated in Figure 3B.

Information becomes normalized in the XIS framework so that any application can use an object's information by getting attributes and relationships dynamically at runtime. Therefore, instead of hard-coding a visualization program to a specific type of object, the visualization looks for the attributes it needs to perform its function.

15 Thus, the developer of the data consuming module may obtain the information needed from the XIS LeifDataItem class, which wraps raw data item objects and exposes them to the data consumer through uniform APIs. The visualization can access all the information exposed, through the XIS APIs and techniques, as needed. If the data source provides AttributeDescriptor objects or translates the properties to domain
20 attributes, then the information can be interpreted more intelligently in domain-specific ways. In one embodiment of the invention, the XIS framework also offers APIs and a programming model (or design pattern) that extends Sun's JavaBean pattern with XIS information awareness. These extended JavaBeans comprise "smarter JavaBeans" that are provided in the XIS framework as the specially
25 configured INFOBEAN objects and, like JavaBeans, INFOBEANS can be visual or nonvisual components, and can be graphically combined in any standard Java development environment.

InfoModel Class Objects

In one embodiment, there are four types of InfoModels in the XIS framework:

BaseInfoModel:

A single BaseInfoModel is required because it is responsible for creating and
5 caching every BaseDataItem (data item) wrapping a Java object (the raw data item
object). The BaseInfoModel object holds any object used as a LeifDataItem until the
data item is no longer needed (i.e., until it is no longer strongly referenced).

SelectableInfoModel:

A super class of BaseInfoModel, this InfoModel delegates to any other
10 InfoModel (typically the BaseInfoModel), and creates SelectableDataItems (data
items that are selectable). The SelectableInfoModel objects add a selected Boolean
attribute and delegate to the LeifDataItems from the prior InfoModel. Since the
BaseInfoModel is a SelectableInfoModel, all LeifDataItems have a selected attribute.
This is true because any LeifDataItem is either in the BaseInfoModel itself, or is in an
15 InfoModel that is nested in the BaseInfoModel and, therefore, inherits the selected
attribute from the BaseDataItem. Additional SelectableInfoModels can be
instantiated to create independent contexts for selection state, independent of other
InfoModels.

AttributeFactoryInfoModel:

20 This InfoModel allows AttributeFactory classes to be registered so that every
time a LeifDataItem is created, it enables attribute factories to add additional
attributes to each LeifDataItem.

InfoModelSubset:

A super class of SelectableInfoModel and AttributeFactoryInfoModel, this is
25 the simplest InfoModel. It delegates to its parent InfoModel and creates
LeifDataItems that wrap those XIS data items from the other InfoModel. This
InfoModel does not automatically add any attributes to the LeifDataItems it creates

and manages, but simply provides a context in which data items can be managed. Each InfoModel "layer" provides a context in which additional attributes can be defined, or existing ones can be overridden. Figures 34A to 34D list an exemplary Java package that lists various classes. Figures 35A to 35D list the InfoModel interface. An interface is a contract in the form of a collection of method and constant declarations. When a class implements an interface, it promises to implement all of the methods declared in that interface.

Exposing Data

Figure 4 is a diagram showing how information from source components may be exposed to data consumers using the mediation layer 112. The Attribute/Metadata interface 214 of Figure 2 exposes a data item 440, 442 as a set of attributes, each of which carry a value and a set of metadata that describes the value. A data consumer component 420, 422, 424 may obtain information from one or more data items. Similarly, a data item 342 may be used by more than one data consumer 422, 424 via the mediation layer or API 112. The information blocks 460, 461 show the information available to data consumers.

A data source component 102 may comprise a data source interface (DSI) object 412, 414. A DSI object 412, 414 provides a conceptual way to encapsulate objects that provide data to the XIS framework 100. A class of a DSI object does not have to extend or implement any interface. Data is created (instantiated) in DSI objects to thereby encapsulate the data and make it available to other objects. The DSI itself may also be a data item that simply contains other data items (members). The data may be generated internally or extracted from an external source such as a database or across a LAN. In addition to creating the data, DSI objects are also responsible for maintaining and controlling data, such as removing and updating the data themselves if such changes are observed in the underlying data used by the DSI to create the data. Property change events (216 of Figure 2) may be used by DSIs to

communicate changes observed in the original source data items to the listening data consumers.

Any Java object may be a data source component. The Swing JButton, for example, is usually a transient object, and not typically retrieved in an information system. The Swing JButton, however, may still be a data source, if so desired.

INFOBEANs as stated above are data consumer components 420, 422, 424. They are specially configured objects that behave like JavaBeans that process and manipulate data in a generic fashion. They are also capable of interpreting and optionally displaying XIS data items. INFOBEANs use the XIS framework API's to gain access to information about data items.

As previously stated, any Java object may become a data item 440, 442 within the XIS framework 100. To make a data item out of JSlider (javax.swing.JSlider), it simply must be added to an INFOBEAN (data consumer component). Table II below provides programming code that makes a data item out of javax.swing.JSlider, thus exposing the information blocks 460, 461 shown in Figure 4.

Table II	
<pre>Jslider slider = new Jslider(0,100,50); tableInfoBean.addRowDataItem(slider);</pre>	

The first line of code creates a data item out of JSlider. The second line of code adds the JSlider created in line 1 to the TableInfoBean object (an INFOBEAN). In this framework, JSlider knows nothing about the XIS framework, however, all its attributes are exposed, to be used and modified.

There are basic steps that must be performed when creating a DSI. First, before any code is actually written, a determination must be made as to how the original data or data source is to be obtained (e.g., will it come from a database, a flat

file, an EJB (Enterprise JavaBean), read from a socket, etc.). This step is performed prior to and independent of coding the DSI. Once the data source has been determined, it must be decided what attributes of the data source should be exposed within the XIS framework.

5 The DSI component within the mediation layer of the XIS framework has the capability of exposing the attributes of a data source. In one embodiment, assuming that the data source is a Java object, the attributes to be exposed are determined through Java introspection. Introspection is a way to determine a bean's properties, methods, and events. Those skilled in the art will understand that a bean is a reusable
10 software component, which may be combined to build applications. In this scenario, no additional code in the consumer component needs to be written. Although convenient, this approach is limiting, as it does not allow control over which attributes are exposed, and does not allow the user to define custom visualization components or express semantics. Furthermore, it requires the Java object to conform to the
15 JavaBean design pattern.

Another approach to determining which attributes of a data source should be exposed involves writing new code. This approach embeds XIS-aware code directly into the Java object. This allows the component developer to directly specify details such as the attributes (e.g., via AttributeDescriptor) to be exposed along with their
20 metadata, the Domains to which the DSI subscribes, which Domain methods are exposed with what implementations, and more. Writing new code enables greater control over the data item (as opposed to the introspection-only case) and provides more flexibility in dealing with any Java object rather than just JavaBean objects. One limitation of this approach is that the relationships between data items may not be
25 directly specified.

Another approach for exposing attributes is to create a Translator class. An application developer may place any of the XIS-aware code for exposing attributes,

Domains, and the like in this class, as specifying relationships to other data as members or as references. This is described further below, in conjunction with the description of Figures 21 to 23.

Referring back to Figure 4, a data consumer 420, 422, 424 may access a data
5 source via a data source interface (DSI) object defined in the mediation layer 112. Communication between the DSI objects 412, 414 and data consumers 420, 422, 424 in the framework 100 is mediated through a set of application programming interfaces (APIs) 112 that remain independent of the actual contents and types of the data sources. Such APIs are contained in class libraries.

10 Using the framework 100, a data consumer 420, 422, 424 (for example, an INFOBEAN as shown in Figure 6B) may be able to use heterogeneous data sources and, thus, eliminate the need to write a particular display software object for each particular data source. Within this framework, any data consumer object or application may use the data source object's information by getting attributes and
15 relationships dynamically at run time. Instead of hard coding a visualization to a specific type of data source component, the visualization object or data consumer looks for the attributes it needs to perform its function. Thus, in this framework, developers for source and consumer objects may work independently of each other.

The DSI object may create any data structure containing members and
20 references, and may expose any part of that data structure to the framework or system 100 if it is desired. Since there are no requirements involved to be a data source, a visualization bean can also be a data source. The data source can consume data from other data sources, and can create new data based on what it has learned. For example, DSI objects may be coded to use data source from an RDBMS, e.g., systems
25 such as provided by Oracle Corporation of Redwood Shores, California, USA and as provided by Microsoft Corporation of Redmond, Washington, USA through their "ACCESS" application or SQL Server. DSI objects also may use data sources such as

XML format data, tabular data (e.g., spreadsheet data such as spreadsheet files in the format of the "Excel" application from Microsoft Corporation), email, schedules (e.g. schedule data in the format of the "PROJECT" application from Microsoft Corporation), data following the SNMP (simple network management protocol), and
5 the like.

Objects created by DSIs are considered to be raw objects in the XIS platform. These raw objects become normalized objects within the XIS framework when they conform to the mediation layer or the API 112. In one embodiment, when these raw objects are added to an XIS-enabled environment, the XIS framework wraps them
10 with a normalized object called a LeifDataItem (see Figure 3A) and holds the data item in one or more InfoModels. The raw data item could be a simple Java object (that does not know anything about XIS) or it can support XIS with direct references to AttributeDescriptors and other metadata that are recognized by XIS. The mediation layer is implemented in software and runs in-process. It is also exportable
15 to serve as a distributed middleware if required by a given application.

An attribute is any property of an object. An attribute 474, 475 typically includes identity (ID number) 476, 477, name of the attribute 478, 479, value 480, 481, and the metadata type or Type Metadata 482, 484, and a description 484, 485. Attribute names 478, 479 may be the table column header or field name in an
20 RDBMS. The description of the attribute 484, 485 may be plain text that explains the attribute. The value 480, 481 may be expressed as a number, string, or other basic data type, or may be expressed as a complex, structured object in itself.

Referring back to Figure 2, the Attribute/Metadata interface 214 provides for fine-grained exposure of an information element or data item 440 442 (Figure 4); by
25 "fine-grained" is meant that the data item is broken down into a number of aspects or characteristics (the "attributes") which are simpler data elements. These might be numerical values, strings, arrays, or anything else that can be represented by a

software object (but is simpler than the original element). If a data consumer 412, 414 is not able to handle the entire data element, it may be able to handle some of its attributes and therefore may be able to do something useful with it, unlike most systems providing interfaces for data integration, which require the entire interface to
5 be implemented to enable any data usage.

The Attribute/Metadata interface 214 also makes each attribute self-describing, to a limited extent. In accordance with the invention, a data consumer designed without regard to a particular type of attribute is still able to display and perform limited operations with those attributes by using the metadata types. For example, a
10 data consumer object or INFOBEAN, which displays an X and Y chart of cost of living (expressed as a numerical value) versus geographical location (expressed as a territory or state character string, e.g., "California"), may be used to display a different set of data contained in another data source component or data item, so long as the metadata types of the different sets of data are compatible with the X and Y chart. In
15 particular, a data consumer component that renders and edits information may still be executed, the default value substituted in cases of absence, and statistical summarization performed, even if the data consumer component is processing a different data source component. Other operations are possible and additional metadata may be provided by developers for any identified classes of information
20 elements (such as number, array, geographic entity, etc.) desired. Because of this, data consumer developers may identify the type of auxiliary information or procedures provided by the consumer component and then define the interface according to which other developers creating data sources should provide or expose their metadata.

25 In one embodiment, any property or attribute of an object may be defined with an AttributeDescriptor programming code. When a data source object is enabled within the XIS framework, the computing system automatically creates

- AttributeDescriptors for JavaBean properties and for some public properties. A JSlider object, for example, is a JavaBean. JavaBeans have properties that are defined by the get<property> and set<property> methods. If the object is a JavaBean and provides a BeanInfo, then it will be honored and treated according to the JavaBean's specification for BeanInfo classes. Only properties listed in the BeanInfo appear as attributes. If metadata types, further discussed below, support the property's type, then they appear as attributes for that data item. Any properties that do not have TypeMetaData support become data items themselves (using the same reflection and introspection rules) and are listed in the data item's reference array, retrieved by the
- 10 LeifDataItem getReferences() method (i.e., a method to obtain the references):
- Translators translate various data item classes, thereby enabling seamless integration of various data items without code modification. Translators may also be added to augment object reflection and introspection, i.e., by developers directly specifying the properties of an object.
- 15 Table III below lists exemplary programming code to define attributes. Table III shows how to define a SIZE attribute.

Table III
<pre> public static AttributeDescriptor SIZE; static { AttributeDescriptorFactory factory = AttributeDescriptorFactory.getAttributeDescriptorFactory(); SIZE = factory.createAttributeDescriptor("size", Your.class, new NumericTypeMetaData("Size", long.class, UnitsOfMeasure.BYTES, 6)); </pre>

}

Metadata are data or information regarding data, for example, data type, field name, length, value restriction, color, and the like. In practice, metadata are used to
5 define the structure and meaning of data objects in tools, databases, applications, and other information processes. The data type metadata 482, 483 encompasses both procedural and declarative information, including (as appropriate for the value) but not limited to the following:

- 10 - unit of measure (from the provided unit of measure software object, converters to/from other units are accessible);
- content length;
- default or maximum character length or precision of numbers;
- data quality (a rating given by the provider of the data);
- 15 - default value (given by the provider of the data type);
- constraints (e.g., numeric range or list of acceptable values);
- formatting;
- comparator function (a procedure which takes two items of the data type in question and returns whether one is to be ordered
- 20 before, after, or the same as the other);
- validation function (returns 'true' or 'false' depending on whether the value for the attribute passes a certain test or set of tests, tests are developer-definable based on an interface which takes a software object providing a value and returns a Boolean value,
- 25 but several reference implementations are provided, such as one to determine whether a number is within a specified range;

in the property sheet (and the display button "Apply" is pressed), the chart automatically updates. Similarly, right clicking on the part of the chart corresponding to a data item enables a user to bring up its property sheet. In this embodiment, this kind of coordination is maintained automatically within the XIS framework through the use of the JavaBeans event mechanisms. It should be noted that when the application is run or executed, the color of the chart does not match the one set in the property sheet. This is because, although the PropertySheet recognizes that any field of type "java.awt.Color" is displayable and editable as a color, the chart has no way of knowing which color field (if there are more than one) should be used. This may be solved, however, within the framework by using standardized Domain Attributes.

In one embodiment of the invention, Translators are used within the XIS framework. As described earlier, a DSI exposes data information within the XIS framework for data consumers to use. Java reflection and introspection may be used to obtain attributes specified in simple objects and JavaBean, respectively. The XIS framework first searches for a Translator class for the data source object. Translators are therefore optional considering that the attribute values and attribute descriptors may directly be obtained from an object. Translators may also be used in order to keep the object's class simple so that it is capable of being used in non-XIS environments.

In an embodiment, Translators are somewhat similar to the BeanInfo class in JavaBeans. Like the BeanInfo class, the name of the class must begin with the name of the data class followed by "Translator". For example, a DSI class named PlanObject.java would have a corresponding Translator class PlanObjectTranslator.java. Also, similar to BeanInfo classes, Translators must reside in the same class space in which the DSI is located. For special circumstance cases such as non-localized objects or custom Translator definitions, a Translator registry,

e.g., TranslatorRegistry, may be used to register a Translator for a given data item class.

Table VIII below shows an exemplary code to register a Translator within the XIS framework.

5

Table VIII
<pre>// register the translator for File objects TranslatorRegistry.getTranslatorRegistry() .registerObjectSchema(PlanObject.class, PlanObjectTranslator.class);</pre>

The above code, as well as all code discussed herein, is provided to illustrate the functional behavior described herein. One of ordinary skill in the art will be able to produce the code, for example, to create classes, objects, APIs to affect this Translator registry.

When the value of an attribute is requested from a data source component, using the mediation layer, the value from the attribute methods specified in the Translator, if any, is exposed. If the Translator does not exist, then the method in the object is invoked, i.e., the get+attribute method within the data source object. The Translator may directly invoke the object's methods as well. The Translator may also be used to wrap legacy code. For example, if the object class is a legacy code, which may not be modified, and has a color property, a Translator may be written to return the value of the color, for example, for the DisplayDomain.color attribute. In this way, the information may be accessed without modifying the original legacy code. The data in the data source object in effect has been normalized or tailored to the

domain policy specified. The XIS framework in this embodiment provides the ability to integrate existing data item classes without modifying them.

To integrate a data item class into the XIS framework, an accompanying Translator class should be provided. In this Translator class, the data source
5 developer provides code that reaches back into the data item's class to expose its properties. It also serves as a way of normalizing access to these properties by imposing a common interface for all Translators to implement.

Translators have two primary functions:

- (1) Translators provide an intermediate method from the caller of an attribute
10 "get" or "set" method to the data source object, enabling a developer to translate the value from its internal representation to a representation understood within the framework, particularly, by the data consumer components. For example, an object may store a default color as a string color name, but the Translator may convert that color to a java.awt.Color object.
- 15 (2) Translators also enable a developer to define attribute types more precisely (e.g., via AttributeDescriptor code, which provides detailed TypeMetaData and may map attributes to specific domains). This resolves any ambiguity in the attribute type. It should be noted that reflected properties that are converted into XIS attributes carry minimal semantics, which contributes to the usefulness of the XIS framework. For
20 example, "float getDegrees()" by itself does not imply whether it provides an angular measure or a temperature measure, nor does it imply that the valid range is from 0 to 360 or from 0.0 to 2.0 pi for angular measures, or distinguish between Fahrenheit, Celsius, or Kelvin for temperature measures.

Translators also enable developers to convert objects that do not follow the
25 JavaBeans get/set API design pattern. For instance, a data item could be implemented as a mapping of attribute names to values (e.g., String, Integer, Double, and Boolean). The Translator may map actual XIS attributes to these key/value combinations,

making XIS data items appear as if they were generated from the more typical JavaBeans-style objects. Example 3 shows how a Translator may be implemented in the XIS framework.

Example 3

5 Example 3 application is contained in three source files listed in Figures 21, 22A, 22B, 22C, 23A, 23B, 23C, and 23D. The resulting outputs are shown in Figures 24A and 24B.

 Example 3 functions similarly to Example 2 above, but enable finer control over the data item attributes exposed within XIS. In particular, the HelloWorld object
10 is explicitly wrapped in a Translator rather than simply feeding it to XIS and letting it rely on reflection to display and manipulate the data. This allows some properties of the raw data item to be hidden, and allows other ones to be better "understood" within XIS. The HelloWorldTranslator.java file contains the Translator.

 Referring to Figure 21, the property change event-handling portion of code
15 1506 found in Figure 18A has been removed from the set methods. These functions are now handled in the HelloWorldTranslator class (Figures 22A-22C), which is registered with XIS by the TranslatorRegistry static call at the beginning of that class.

 Referring to Figures 22A to 22C, this file contains the HelloWorldTranslator
20 class 2202 that wraps the HelloWorld object. When an INFOBEAN encounters a raw data item, it ultimately accesses the raw data item through a Translator's methods, if a Translator exists. In one embodiment, all Translators are subclasses of the com.xis.leif.im.Translator class, which provides an infrastructure that facilitates the use of attributes already defined or set-up in an XIS Domain. Such attributes have type metadata predefined for them, and many XIS INFOBEANS automatically access
25 and utilize the metadata and the attributes themselves.

 The first member variable 2204 of the HelloWorldTranslator class lists the predefined XIS domains from which attributes are taken. In order to use an attribute

not in one of these domains, it may be necessary to provide a separate metadata for it. The second member variable 2203 defines an array, which lists auxiliary information on the fields, in this case, whether they are displayed by default in the property sheet display. The getFieldMetadataArray method 2206A and 2206B fills this array, if
5 necessary, with the default metadata for the fields that are exposed by the get and set methods further listed below in the source file. The exception to the defaults is that the "course" attribute 2208 is not visible, i.e., it is not displayed under the "Preferred Attributes" on the property sheet. It is displayed, however, when the "All Attributes" option is selected. Other things that may be set (but are not in the source file) include
10 the sorting order and rank when multiple data items of this class are listed together, as in a table.

The ensuing get and set methods 2210 define the attributes that the rest of the, XIS objects can see. In general, this refers to the enclosing raw data object, but this need not always be the case. For example, here, the HelloWorld's integer 'value'
15 property is map to the 'speed' attribute, which takes its metadata and its type (double) from the Movement domain 2212. The Movement domain has previously been defined. By handling conversions within the Translator, the rest of XIS "never knows" that there is really an integer value underneath.

All change event propagation is handled through the request objects passed
20 into the set methods 2214. Similar to Example 2, all XIS objects using the same data item are notified of the change.

Referring to Figures 23A through 23D, the TestHarness code is largely unchanged from Example 2, but at the end it includes a separate class that demonstrates how attributes may be altered on a data item that includes a Translator.
25 In general, the procedure is to obtain, first, the LeifDataItem corresponding to the raw data item, then a "domain wrapper" that wraps this and enables access to attributes that come under that domain. Internally, the domain wrapper calls methods on the

Translator, but externally it presents an interface to the developer that depends only on the domain and not on the particular LeifDataItem being wrapped.

The code in the Accelerate class is a standard thread implementation that uses the wrapper to successively change the HelloWorld data item's speed attribute. In
5 other embodiments, this attribute may be altered by any portion of the program that has access to the data item. Because the wrapper eventually calls the Translator, the change events are propagated appropriately.

The Example 3 application, similar to Example 2, displays two windows--a property sheet and a chart. Example 3 shows that the property sheet may be made to
10 display more properties (in this case, the "course" attribute) by selecting "All Attributes" from the top menu, which ignores the "preferred" setting in the FieldMetaData. The attributes all have reasonable names, which come from the metadata stored in the domains. Finally, if the "Pen Color" attribute is edited, the plot changes color accordingly. This is because the Chart INFOBEAN recognizes and
15 uses the "Pen Color" attribute from the Display Domain if it finds one on the data item. Because of the various wrappers employed, it does not matter what kind of object is eventually underneath or what other attributes it has; as long as the "Pen Color" attribute is found, it will be used.

Example 4

20 Example 4 contains two source files namely HelloWorld and TestHarness. Figures 25A to 25C list the HelloWorld.java source file. This example shows translation without a Translator class. The resulting output is a property sheet INFOBEAN and a Chart INFOBEAN. Figure 26 shows the property sheet InfoBean, but the ChartINFOBEAN is not shown. The HelloWorldTranslator.java file from
25 Example 3 is eliminated. Similar to above, the HelloWorld.java file has been marked such that changes from Example 3 are marked by open and close commented braces.

The Java source files enable some of the same fine control over the attributes exposed within XIS (see Example 3) without using a separate Translator class. If a developer has control over the Java code for the data source, the overhead of using a Translator may be eliminated by building some of its functions directly into the data source. It may, however, still be desirable to use a Translator because this avoids building any XIS-specific code into the JavaBean created and enables the code to be more streamlined for use in non-XIS contexts.

Figures 22A to 22C contain the HelloWorld code, which include some of the code from Example 3, with some additional code for exposing attributes and providing field metadata as the HelloWorldTranslator class in Example 3. In fact, almost everything may be done from within this data item class that could be done with a Translator except for two things. First, attributes cannot be hidden from XIS--everything with a get() method is exposed. Second, Translators may be used in conjunction with security features, for example, embodied in the XIS framework to restrict access to visible properties in a generalized fashion across all data items.

The first portion of code 2502 provides methods for XIS to obtain information on which data item properties have attribute information provided for them. The XIS components first introspect on the data item to find its properties, then try to call the get(property name)Descriptor for each property. For those instances where this fails, the XIS component may provide its own attribute descriptors where it recognizes the type of the attribute. For example, the property sheet recognizes Color, Numeric (various subtypes), Date, and String attributes. The next portion of code 2504 initializes a PropertyChangeSupport instance, because property change events must be handled here, similar to Example 2.

Following this are several methods unchanged from Example 3, but the set() methods have been altered to resemble those from Example 2 in which property change events are fired. Next, there are several new get() and set() methods which

expose the properties with the same names as the Translator did in Example 3. These names correspond with the `getAttributeDescriptor` methods listed above in the code, and the XIS components interpret them accordingly. These properties show up as preferred attributes in the property sheet, whereas those properties without
5 corresponding attribute descriptors become nonpreferred attributes.

The portion of new code 2506 following this provides field metadata for this raw data item. This program code is almost identical with those in the Example 3 `HelloWorldTranslator`, except that both the variable and the method are now static, and the "Course" attribute removed to reduce clutter.

10 The TestHarness file for Example 4 is generally unchanged from Example 3, except the code to register a Translator is removed. The portion of program code 2302 in Figure 23A is removed in this example.

This Example 4 application is largely similar to Example 3, except that if the "All Attributes" option is chosen under the property sheet, then the previously
15 unexposed "value", "myColor", and other attributes appear. The chart is also able to display the "value" and "ID" properties.

Figure 26 shows a window display of the property sheet output from the Example 4 application.

Processing Flow for Exposing Data

20 Figure 27A and Figure 27B show a flow diagram of how information about a data item, e.g., a data item as shown in Figure 5, is exposed to data consumer components. In the preferred embodiment, data exposure within the system may be accomplished by one of three methods, described further below. There are two primary functions to be performed: (1) interfacing with an internal or external data
25 source such as a set of disk files, a database, or a web service; and (2) exposing the resulting data according to the standard interface outlined above, e.g., following the domain interface, relationships interface, and the like.

In the first method of exposure, one software component performs both functions; in the second method, separate software objects perform the two functions; in the third method, which is only implementable in an object oriented programming language that provides self-analysis facilities, such as Objective C or Java, the exposure function is performed automatically by the system. More particularly, in the first method, a data source interface (DSI) software object is written that connects to an external data source such as a local file system, an Internet site, or an RDBMS. The DSI object converts the information provided by the external data source into the information representation or data item that may be consumed by data consumer objects as described above. In a preferred embodiment of the invention, the data source DSI objects are from the LeifDataItem class.

In the second method, the API of an existing software data object is accessed by a Translator object, which then exposes the information to other data consumers in the representation format discussed above. This arrangement is referred to in the object oriented design literature as an Adapter or Wrapper pattern. See, for example, Gamma, E., Helm, R., Johnson, R., Vlissides, J. (1995); Design Patterns: Elements of Reusable Object Oriented Software; New York: Addison-Wesley. This is generally referred to as "wrapping". More precisely, a data consumer component requests information from the Translator object, which then converts the request into something the original data source object can deal with. The consumer then converts that object's response into the appropriate format for a response to the original call.

The third method applies the facilities within a self-analysis-capable programming language such as Java or Objective C. This type of programming language enables an object to determine the available data fields and function calls afforded by another object while an application is running. In the preferred embodiment implemented in Java, the queried object must be a "JavaBean" according to Sun Microsystem's suggested convention, with "get()" and "set()" method calls.

Each data field available through one of these calls, or by virtue of being a "public" variable, is examined, and the resulting type information is used to expose both the field's value and appropriate metadata to data consumers. If the field is a numeric type (double, Double, float, Float, Integer, int, long, or byte in Java), it is exposed as a decimal or integer with an "unknown" unit of measure. If the field is a string (String in Java), it is exposed as a text element. If the field is a complex object, it is exposed as a linked element via a reference (discussed above), unless a special class has registered to handle the complex object type within the application context or XIS framework.

In another embodiment, in addition to the facilities for exposing data, the XIS framework provides an apparatus for allowing application context to affect how information is exposed. This is illustrated in Figures 27A, 27B, and 28, described further below. The attributes, methods, and relationships available for a given data source are determined by a surrounding structure called an InfoModel. The InfoModel connects the pieces of the XIS framework together. It is an interface comprising an object that knows how to return data item information. See also Figure 3A discussed above.

In another preferred embodiment, an application comprises one or more Views, each containing and providing an InfoModel to one or more data consumers, termed INFOBEANS as described above, which may provide display, computational, or interactive facilities. All of the Views are in turn managed under a single ViewHost, which provides the global application context. Conceptually, a view represents a single perspective of the current data. One single perspective (or view) could be quite different from another because of the ability to add/override attributes in a view. View creates a new InfoModel that in turn creates new LeifDataItems that are used within that view. This is done to separate context sensitive attributes between an InfoModel and another View, INFOBEAN, or InfoModel. Views are also used as a

controller of one or more INFOBEANS by managing the data flow to each individual INFOBEAN. All Views are also by definition INFOBEANS since they manage information flow with generic means. The ViewHost is used to manage the complete set of Views and INFOBEANS, thus presenting each view to the user in a uniform
5 fashion such as frames, windows or other display mechanisms.

Figure 27A and Figure 27B show a flow chart of a data exposure facility within the XIS framework constructed in accordance with the present invention. The diagrammed operations will occur when the XIS framework is loaded into program memory of a computer system that supports an object oriented programming (OOP)
10 environment, as will be understood by those skilled in the art. In the first processing step, at block 2702, a data source object is registered within the XIS framework of the computer system. When a data item is brought into the framework, it is exposed to data consumer objects, e.g., via an InfoModel, typically via one of three methods. The first operation is to determine whether a Translator object is available and whether the
15 Translator object is registered for the raw data object type. This processing is represented by the decision block 2704.

If the Translator is available, an affirmative outcome at the block 2704, then at block 2760 the data item attributes, methods, and Domains are obtained via the Translator. If the Translator is not available, a negative outcome at 2704, then at
20 block 2706 the data item object is scanned to find the data object's methods. The scanning may be accomplished with facilities of the OOP environment of the host computer system, such as InfoBeans.

After the data item object has been scanned at block 2706, the system next checks to see if the data item includes an XIS-standard data exposure interface (that
25 is, a predefined interface known to the XIS framework). If the standard interface is available, an affirmative outcome at block 2714, then at block 2762 the data item object attributes, methods, and Domains are obtained from the data source object

itself. The system then checks for references to Domain Policies at block 2763. If there is no standard interface available at the decision block 2714, then processing moves to the block 2764, where the facilities of the programming language of the OOP environment are used to interrogate the data item object to determine the object's accessible data fields. Processing then moves to exposure attribute processing via the off-page connector B to Figure 27B, described below.

Returning to block 2762, after the data item attributes, methods, and Domains are obtained, a check is made for Domain Policy references at the decision block 2763. If there are references to Domain Policies, an affirmative outcome at the decision block, then at block 2765 a check is made to determine if there is FieldMetaData override for a Domain. If there are no Domain overrides, then at block 2710 the definitions from the referenced Domain Policy are used for the attribute metadata. The processing then moves to exposure attribute processing via the off-page connector A. If there are Domain overrides, an affirmative outcome at block 2765, then at block 2716 the system uses the definitions included with the data item object Translator for the attribute metadata of the data item. The processing then moves to exposure attribute processing via the off-page connector A. Referring back to block 2763, if there were no references to Domain Policies, then processing proceeds straight to block 2716 to obtain the attribute metadata definitions and then to the exposure attribute processing via the off-page connector A to Figure 27B.

Turning to the processing of Figure 27B via the off-page connector A, the data item attributes are exposed to the local InfoModel class (i.e., a DSI object). Next, at the decision block 2770, a check is made to determine if the data item has any dynamic attributes. If the data item does, then at block 2772 the framework processing adds the dynamic attributes to the available data item attributes. If the data item has no dynamic attributes, or after any dynamic attributes have been added, processing moves to block 2773, where the system checks to determine if a Translator

was provided (registered) for the Data Source object. If a Translator was provided, then at the decision box 2775 the system checks to determine if reflection has been requested. Those skilled in the art will understand that reflection is an operation that is provided by most OOP environments, such as by Java.

5 If reflection is to be performed, then processing returns to Figure 27A via the off-page connector C, whereupon the Data Source is scanned. If reflection is not to be performed, or if no Translator was registered (a negative outcome at block 2773), then processing moves to the decision block 2774. At block 2774, the system checks to determine if the data item has registered any extended Translators. If there are no
10 extended Translators, then at block 2718 the data item attributes are derived, added, or hidden, depending on the data context gleaned from the DSI. The data item is then fully exposed to data consumer components for use (block 2790). If there are extended Translators, an affirmative outcome at block 2774, then processing returns to Figure 27A via the off-page connector B to obtain the information from such
15 Translators (block 2760).

Figure 28 is a diagram showing how InfoModels provide contextualization within the XIS framework. The initial "raw" object (e.g., Java object or data from RDBMS) provided by the data source is exposed through mediation by successive "wrapping layers" of other objects, which determine which fields are visible, and
20 possibly performs conversions or consolidations. Each "View" within the XIS framework, which contains one or more "INFOBEAN" consumers, carries with it its own context. Feeding a data source component to a consumer component places it within that context.

When a data element provided by a data source is introduced to an InfoModel,
25 the InfoModel actually provides a new element derived from this one to consumers. This derived element provides the actual data source interface discussed above. The attributes, type metadata, and methods provided for the element are first determined

as described below under "data exposure facilities". Then the InfoModel inspects the results and decides whether to add or derive any new attributes, or to hide existing attributes. Once the data is added to an InfoModel (discussed below), it may be asked for by other components for processing and visualization purposes.

5 In another embodiment, the ViewHost, a data consumer component, provides facilities allowing interactive cut and paste and drag and drop of information elements between different consumer components. When a cut-and-paste or drag-and-drop occurs, each information element is decontextualized and then recontextualized inside the new InfoModel into which it is placed. This means that some attributes
10 previously available may be removed, whereas others, more appropriate to the functions of the recipient consumer(s), may be added.

Referring back to Figure 28, this figure shows how multiple InfoModels work to share references to raw objects. This figure also illustrates how an INFOBEAN and/or a DSI adds objects to the InfoModel to obtain data items. The DSI 2815
15 creates an instance of source raw data item (e.g., Java object) 2830. The original source objects are maintained by DSI and one or more InfoModels have references to these source objects. The raw data item 2830 is provided to or may be used by an INFOBEAN (data consumer) via the addRawDataItem() method. (Please refer to Table III above. The second line of code shows that the source object is added to the
20 TableInfoClass.)

The INFOBEAN 2820 requests the data item for the raw object 2835 (which was the source object 2830) from its InfoModel or the BaseInfoModel. The BaseInfoModel as discussed above is responsible for creating and caching every data item (wrapping any object). Each InfoModel is considered to contain a subset of data
25 items from the entire set of data. The InfoModel is a view of the BaseInfoModel, thus, there may be data items in the BaseInfoModel that are not present in the InfoModel used by the INFOBEAN. The circles 2828 shown in the InfoModel 2805

represent attributes, domains, relationships, etc. from the BaseInfoModel. For example, Figure 6 contains a list of information that an INFOBEAN may obtain about the Person object, in this case 2830, 2835, through the LeifDataItem API. In addition information from nested InfoModels are inherited. An InfoModel 2805 also always
5 adds new items to the BaseInfoModel, e.g., common attributes such as "selection" are delegated. These new items are also exposed to the INFOBEAN as shown by the number of circles in 2840.

The BaseInfoModel 2810 returns the raw data object 2845 to the InfoModel. The InfoModel 2805 returns a data item 2840, which exposes the information, for
10 example, shown in Figures 4 and 5 to the INFOBEAN 2820. A data item 2845 (see Figure 4) is returned, which may have been created or may have already existed in the InfoModel.

In the XIS framework, as shown in the InfoModel, attribute and method "requests" mechanisms are available. Standard mechanism to retrieve the current
15 value of any Attribute, or set the current value, if the Attribute is mutable, may also be easily implemented. In one embodiment, this mechanism also enables a standard syntax and mechanism for the invocation of any domain method.

These attribute and method "request" objects contain all of the information necessary for the data source to fulfill the requested task. It contains the actual
20 ("raw") data item whose attribute value is being sought or set (along with the desired new value), or whose domain method is being invoked. They also contain contextual information, such as the current User, allowing data source developers to implement full user-based information access and security controls, and a context for obtaining services (implemented through the Java "BeanContext" APIs). This is significant
25 because it enables the process of executing methods and/or getting or setting values to leverage services "in context" such as the current data view (or window) through which to interact with the user. These request objects also contain contextual "data

item" (e.g., LeifDataItem), as opposed to the "raw" non-contextual data item, from which to get additional or overridden values or execute additional/overridden methods, if appropriate to the task. In another embodiment, these "Request" objects are not only passed when getting/setting values, or executing domain methods, but are
5 also passed to all other "Translator" methods that take and may need to access or manipulate the data item. The "Request" object not only provides context, but also ensures that the DSI (via the Translator) always has what it needs (particularly, the User information) to manage its security constraints, if any.

A data item locking mechanism for obtaining, holding, and releasing a lock on
10 an individual data item may be implemented in accordance with the present invention. This implementation is best defined by the data source provider, by implementing such locking mechanism in the data source, itself. Ad hoc implementations for data sources that do not provide native locking may also be done via the DSIs. This mechanism enables DSIs to implement multiple changes atomically (as one
15 transaction) by effecting those changes at the time of the release. In addition, a "rollback" feature for atomic changes is possible through the "revert" capability, which generally releases the lock without affecting the previous changes. This locking mechanism may also be leveraged to lock multiple objects simultaneously, at the discretion of the DSI, by associating the lock with a DSI-defined "lock object"
20 which can either be associated with a single data item or with a group of data items, as appropriate. Using the DSI API, this provides to data consumers a single API for dealing with locking implementations.

Figure 29 shows an embodiment of a "plug-in manager" called PlugInManager of the present invention. The programming framework embodied in this invention
25 includes a "plug-in manager" which provides extensible communication between services and components within an application. Plug-ins are software modules that add a specific feature or service to a larger system (i.e., plug into a larger application

to provide functionality). For example, there are a number of plug-ins for the NETSCAPE NAVIGATOR browser that enables the browser to display different types of audio or video messages, e.g., SHOCKWAVE by MACROMEDIA.

A "service" is capable of obtaining information on a data consumer or data source component and capable of using that information to provide specific facilities to the application as a whole. Examples of such facilities include addition of entries to a top level application menu, loading MIME types into the global MIME type map, and running of scripts upon loading or unloading of data sources or consumers. The plug in manager loads new services at initialization and also at any time through an explicit request. The service in turn queries all modules present in the application as to whether they request any actions of the given service, executing them if so indicated.

In more detail, as each new module, object, or component is loaded into the XIS framework, the PlugInManager 2905 on startup loads a special service called PluggableServiceFinder 2925, which searches for other services 2915 and loads the services found into the PlugInManager 2920. The PlugInManager holds or keeps track of all services. These services are referred to as PluggableServices. As each pluggable service is loaded, the PlugInManager searches such service for defined resources and attaches these resources to the appropriate components or performs actions on resources found 2935. For example, a menu manager service might search for resources specifying menu entries and associated operations. A component that wants to add a menu entry to the global application menu bar would simply provide a resource of this type. In essence, as each new module is loaded, the PlugInManager tells the currently registered services ("PluggableServices") to look for new plug-ins. Thus, in essence the PlugInManager acts as a plug-in for plug-in handlers.

The combination of the PlugInManager and the PluggableService interface allows communication between an XIS application and various XIS components. The

components may contain a number of resources that affect certain areas of XIS. For example, they may have a resource file that contains an ECMA script. ECMA is an international industry association dedicated to the standardization of information and communication systems (see www.ecma.ch). If there is a PluggableService that

5 knows about that resource file, then it may decide to run the ECMA script when it finds the resource. The PlugInManager is responsible for managing the resource list and all PluggableServices. When a new PluggableService is added, it is given all the resources to determine if it is interested in any of the resources on the list. If it is interested in a resource, then it may perform some action.

10 In one embodiment of the invention, an ASCII file is created to indicate that a JAR file contains a resource to be provided to all PluggableServices.

Table IX below shows an exemplary ASCII file that indicates that there are resources associated with the module containing the TableView class. It does not indicate what resources exist; it is up to the PluggableServices to determine if there is

15 a resource in which they are interested. The PlugInManager, using Class.forName(), creates the TableView.class and calls the loadPlugIn method on each of the PluggableServices. Below is the content of the loadPlugIn method in the MimeTypesPluggableService:

20

Table IX	
<pre> public static final String RESOURCE_NAME = "leifResources/mime.types"; public void loadPlugIn(Class relativeClass) { if (loadedResources == null) { synchronized (this) { if (loadedResources == null) { loadedResources = new HashSet(); } } } } </pre>	

```

    }
    }
}
try {
    String resource = PlugInManager.convertResourceName(
        relativeClass, RESOURCE_NAME);
    ClassLoader loader = relativeClass.getClassLoader();
    Enumeration enum = loader.getResources(resource);
    while (enum.hasMoreElements()) {
        URL resourceURL = (URL)enum.nextElement();
        if (resourceURL != null) {
            if (!loadedResources.contains(resourceURL)) {
                loadedResources.add(resourceURL);
                LeifJAFUtilities.addMimeTypes(
                    baseInfoModel.getOwnedBeanContextChild(), resourceURL);
            }
        }
    }
} catch (IOException e) {
    e.printStackTrace();
}
}

```

The entire arrangement of data sources separated from consumers by a mediating structure, InfoModels providing context to Views, and a ViewHost and
5 plug in manager playing set central organizing roles, allows a multi-purpose

application to be rapidly assembled out of preexisting components. Each available data source and consumer component may be integrated into the application by the plug in manager, and compartmentalization.

Figure 30 shows a flow diagram of an embodiment of the invention that assists
5 in creating XML (extensible markup language) DSIs and database DSIs. The XIS framework provides two sets of facilities for exposing electronically available
7 information that comes from data sources that are partly structured and/or self-describing, and for writing back possibly modified information to the sources. For example, one of these works for XML formatted data, and the other works for data
10 stored in relational databases. However, the mechanisms are inherently extensible to any other data source that is partly structured and/or self-describing.

The World Wide Web Consortium (W3C), which has standardized most of the XML formats to date, is in the process of defining a new XML format for declaring schemas, called XML Schema. The XML schema addresses the means for defining
15 the structure, content, and semantics of XML documents. The present invention is configured to take into account future versions of the XML schema.

In one embodiment, the present invention includes aspects of the method outlined below that are in common between the XML and relational database interfaces and can be applied in this general capacity to be part of the invention as
20 well. In particular (described further below), an embodiment includes the apparatus involving separate processing of "data" (3002 branch) and "schema" (3004 branch) information, with a type conversion stage 3012 during the former and a metadata creation stage 3014 during the latter, followed by the construction of semantically enriched attributes 3016, 3018 based on the results of both these processes.

25 In one embodiment, one method involves two apparatus or facilities designed to handle extensible markup language (XML) formatted data. An XML text stream comprises of a set of hierarchically structured "tags" interspersed with text "content".

The names, data types, and ordering of the tags may vary from stream to stream and is specified by a separate document called an "XML schema." A description of the XML schema is provided in the document "XML Schema Part 1: Structures (Working Draft)", which may be obtained from the W3C at the WWW address of
5 <http://www.w3.org/TR/xmlschema-1>, and also in the document "XML Schema Part 2: Datatypes (Working Draft)", which may be obtained at the WWW address of
<http://www.w3.org/TR/xmlschema-2>. Given an XML data stream, the first apparatus parses the XML schema appropriate for it and constructs metadata specifications appropriate for the contents of each tag. In some cases, these specifications are in the
10 form of computer program segments that can be compiled into executable procedures that can be accessed immediately by a running application; in other cases, they refer directly to already precompiled procedures. The second apparatus parses the data document, matching the tags up with the types and metadata derived from the schema, and exposes the content in the standardized information representation using this
15 metadata. This information includes the generic relational links described above in addition to attributes and metadata.

In particular, the structures defined by an XML schema document determine the hierarchical and type structure allowed in XML instance documents, and this structure is paralleled again in the information elements (with attributes and member
20 elements) constructed on parsing an instance. The following table (Table X) shows the correspondence:

Table X		
XML Schema	Compliant XML Instance	Resulting Information Element
simple type outside attribute declaration	element with no subelements	(typed) attribute
simple type inside attribute declaration	attribute	(typed) attribute
complex type without descendants	element with attributes but no subelements	information element with attributes
complex type with descendants	element with subelements	information element with member elements

A complexType element, as defined in the W3C XML Schema standard, defines the structure of an information element. Elements contained within the complexType element which are defined as complexTypes themselves, are translated into the member elements of the information element defined by the containing complexType. All other elements and attributes contained within the complexType element, and defining single value types, are translated into attributes of the information element defined by the containing complexType.

The XML ingestion process supports an additional capability that allows for the translation of types defined by an XML Schema, into more semantically rich domain policy based attributes that consumers in the present framework can take advantage of. To do this an embodiment of the present invention includes a process that uses XSLT (XML Stylesheet Language Transformations) to scan an XSDL (XML Schema Definition Language) document, and embed domain specific instructions in a new version of the XSDL document. The embedded instructions

direct the XML processing software to create information elements that contain the semantically richer domain attribute representations. In particular, the tags in the table below may be placed by the XSLT within <appinfo> subelements of elements in the XSDL schema document. When the schema parsing apparatus reads the resulting schema, it sets up the XML processor to perform the mappings as shown when reading an instance document. Because of the nature of the handling specified for <appinfo> elements in the XSDL specification (see the W3C Web site at <http://www.w3.org/XML/Schema>), these will simply be ignored by other XML schema processors that have not been designed to read them. Table XI lists properties that may be used:

Table XI	
Property	Description
domainfield	A Domain Policy attribute to map instances of this XML element or attribute to
typemetadata	A metadata type to map instances of this XML element or attribute
converter	A type converter for transformation of XML data values (strings) into typed values such as floating point numbers or date objects
visible	The initial visibility state with which instances of this XML element or attribute should be exposed within an information element
mutable	Whether or not instances of this XML element or attribute can be edited after being exposed within an information element

Conversely, the system has the ability to construct an XML document and associated XML Schema from a hierarchy of information elements. The process for outputting an XML document and associated XML Schema from a hierarchy of information elements is essentially the reverse of the ingestion process. The resulting
5 schema documents are generated with the special <appinfo> instructions that carry the additional semantic meaning for the attributes, but (see above) these additional instructions are embedded in a way that makes them transparent to other XML document readers that do not recognize them. In this way, an application-to-application information communication session can be set up using HTTP or a similar
10 protocol for synchronous communications, and SMTP or similar for asynchronous communications. The sending application encodes information elements into an XML schema/instance pair using the just described apparatus and sends both to the receiver application. The receiver reconstructs the information elements also using the above apparatus. Thus, any form of structured, annotated information that can be
15 expressed within our framework can be transmitted between applications using a simple document transfer protocol. Another embodiment of the present invention involves an apparatus 3020, which will connect to a relational database system and extract or write information from it using the structured query language (SQL). The user of such an information source needs only to specify the tables and columns to
20 access within the database, as well as any selection criteria to narrow the range of data returned. The apparatus performing this function will automatically convert the data in the database into the appropriate representational format and add metadata and parent-child links as required. In general, the process is similar to that described for the XML data input/output capability, except that SQL statements and conversion
25 mappings specified by the user take the place of the XML schema and document type and hierarchy information.

The process translates the result of a database query into an XIS data item, and each row in the result set is translated into a member data item of the result data item. Each row data item can also have a database query associated with it which can be executed to provide its own member data item. Access to the member items is on-
5 demand (e.g., as called for by a consumer), an operation we refer to as "drill-down".

The row data items are created with attributes that are derived from the columns returned by the result set. The attribute specifications can be generated automatically from the result set column SQL type specifications (integer, date, text, etc.), or can be customized to suit the needs of the user. The customization process
10 allows the user to map database result column type definitions into more semantically rich domain policy and/or metadata-enhanced attributes that consumers can take advantage of. In other words, the DSI can associate column data types with metadata types and domain policy attribute types. To do this the framework supports a number of attribute customization properties that the user can specify (described further
15 below).

Modifications to row data items are handled by allowing the user to associate parameterized update, insert and delete queries with a given result set. When the value of a record member attribute is changed, the user specified parameterized query is executed, and the parameters are filled in from the attribute values identified in the
20 query specification. Insertion and deletion of record members are handled in a likewise manner.

The following table, Table XII, describes the list of customizable properties supported by the SQL database interface process. These properties are used to describe a single interface to a database for producing one or more hierarchically
25 organized information elements from its contents.

Table XII	
Property	Description
jdbc.driver	JDBC vendor driver class name
jdbc.url	JDBC database connection URL
display.name	Default display name for result data item
select.query	SQL select statement for result data item
update.query	Parameterized SQL update statement
update.columns	Result column names that fill update query parameters
insert.query	Parameterized SQL insert statement
insert.columns	Result column names that fill insert query parameters
delete.query	Parameterized SQL delete statement
delete.columns	Result column names that fill delete query parameters
drilldown.query	Parameterized SQL select statement for record drill-down
drilldown.columns	Result column names that fill drill-down query parameters
result.attributes	Result data item attribute names
record.attributes	Record data item attribute names

- The names of attributes listed for both result and record data items are used to uniquely identify the attributes. The customization of each attribute is done according to its unique name. The following table (Table XIII) describes the list of customizable properties that can be applied to each attribute, where <name> is the unique attribute name. Note the parallel between these properties and those given for the '<appinfo>' tag insertion in the case above of XML data parsing.

Table XIII	
Property	Description
<name>.label	The attribute display label
<name>.domainfield	A XIS Domain Policy class field to map this attribute to
<name>.typemetadata	A XIS metadata class name to map this attribute to
<name>.converter	A XIS type converter for transformation of result values
<name>.columns	The result columns names that will be converted
<name>.visible	The attribute initial visibility state
<name>.mutable	The editability of the attribute

Given the metadata from known industry standard APIs (e.g., JDBC and ODBC) that provide metadata for results (typically handled as tabular columns) including string, length, numeric precision, date formatting, etc, relational databases may have dynamic DSI toolkit built and implemented as shown in Figure 27. Tabular data may also be interpreted as a default string typed. Java objects, remote Java objects (via RMI or CORBA-derived interfaces), and Enterprise JavaBeans (a variant of RMI objects) may have DSIs automatically created for them, using, for example, reflection and known types (primitives and other typical types like String, Data, etc., and registered types). A developer may also place XIS-specific annotations (e.g., static AttributeDescriptor and FieldData methods) on reflected Java objects. Collection subclassed objects may also indicate containment.

Dynamic data supplemental metadata may also be used in the XIS framework of the present invention. Schemas with little or no expressive information may also be used to specify additional metadata within the XIS framework, which include

mapping attributes to type metadata, mapping attributes or methods to domains, defining new attributes/methods based on existing data from the source, defining functions/user commands for easily implemented capabilities (e.g., deleting a data item, adding/inserting a new data item, initializing a template data item that a user can
 5 fill in before adding, etc.), and defining relationships to other data items to perform "drill-down" queries. Such type of schemas may be obtained, for example for Java Metadata API, CORBA API (on which the Java API was based), Oracle "common warehouse metamodel," and the like. Other formats for mapping metadata to schemas from other sources, or as might be defined in the future, may easily be incorporated
 10 within the framework.

Distributed Application Capabilities

In an embodiment, the application development framework being discussed offers three sets of facilities to aid in the development of distributed applications in which users, software application components, or both are distributed across multiple
 15 computers connected by a network.

CORBA Exposure

The first set of facilities allows the automatic exposure of a data source over a network to any remote consumer software object via the Common Object Request Broker Architecture (CORBA) protocol. This protocol allows the remote consumer to
 20 utilize information elements just as if they existed locally, but usually with lower bandwidth than it would take to send the entire elements over the network connection.

To better understand the invention, examples are shown below in Table XIV, which shows an exemplary XML document.

Table XIV	
<pre><tracks> <track name="FIRST"></pre>	


```
<country>US</country>
<category>SUB</category>
<threat>HOS</threat>
<position>
  <lat>10.0</lat>
  <lon>10.0</lon>
  <alt>10.0</alt>
</position>
</track>
<track name="SECOND">
  <country>US</country>
  <category>AIR</category>
  <threat>FRI</threat>
  <position>
    <lat>20.0</lat>
    <lon>20.0</lon>
    <alt>20.0</alt>
  </position>
</track>
</tracks>
```

The following table (Table XV) shows the result if the above XML document were run through an XML DSI.

5

Table XV
LEIF Data Item - tracks

LEIF Data Item – track

Attribute - name : String value='FIRST'

Attribute - country : String value='US'

Attribute - category : String value='SUB'

Attribute - threat : String value='HOS'

LEIF Data Item - position : LEIF Data Item

Attribute - lat : String value='10.0'

Attribute - lon : String value='10.0'

Attribute - alt : String value='10.0'

LEIF Data Item – track

Attribute - name : String value='SECOND'

Attribute - country : String value='US'

Attribute - category : String value='AIR'

Attribute - threat : String value='FRI'

LEIF Data Item – position

Attribute - lat : String value='20.0'

Attribute - lon : String value='20.0'

Attribute - alt : String value='20.0'

Table XVI below shows an exemplary XSDL document.

Table XVI

```

<?xml version="1.0" encoding="UTF-8"?>
<schema xmlns="http://www.w3.org/1999/XMLSchema">
  <element name="tracks">
    <complexType>
      <element name="track" type="Track" maxOccurs="unbounded"/>
    
```

```
</complexType>
</element>
<complexType name="Track">
  <attribute name="name" type="string"/>
  <element name="country" type="string"/>
  <element name="category" type="string"/>
  <element name="threat" type="string"/>
  <element name="position" type="LatLonAlt"/>
</complexType>
<complexType name="LatLonAlt">
  <element name="lat" type="decimal"/>
  <element name="lon" type="decimal"/>
  <element name="alt" type="decimal"/>
</complexType>
</schema>
```

- XML Schema Definition Language (.xsd) documents are used to express the semantic descriptions of Attributes and the Domains in which they live. Code
- 5 generation utilities as described above create the actual Domain Objects with their groups of related Attributes. Thus, developers may create new Domains, and add Attributes to their existing Domains, by modifying these XML Schema files and generating the appropriate Java code. XML Schema files include basic attribute constraints (e.g., value ranges, list of valid values, attribute naming, formatting, etc.)
- 10 that are turned into TypeMetaData objects, as well as textual comments that are converted into Java documentation to define the intent of the attribute (or method, in

the case of Domain methods). In some cases, Java code is embedded in the document and used inline in the generated Java classes for the Domain.

Table XVII below shows the result when the above XSDL document is used to type the elements in the XML document source. Note that the type for lat, lon, and alt are now Double instead of String.

Table XVII	
LEIF Data Item – tracks	
LEIF Data Item – track	
Attribute – name :	String value='FIRST'
Attribute – country :	String value='US'
Attribute – category :	String value='SUB'
Attribute – threat :	String value='HOS'
LEIF Data Item - position	
Attribute - lat :	Double value='10.0'
Attribute - lon :	Double value='10.0'
Attribute - alt :	Double value='10.0'
LEIF Data Item – track	
Attribute – name :	String value='SECOND'
Attribute – country :	String value='US'
Attribute – category :	String value='AIR'
Attribute – threat :	String value='FRI'
LEIF Data Item - position	
Attribute - lat :	Double value='20.0'
Attribute - lon :	Double value='20.0'
Attribute - alt :	Double value='20.0'

The following table (Table XVIII) shows an exemplary XSLT document.

Table XVIII
<pre> <xsl:for-each select="xsd:element"> <xsl:choose> <xsl:when test="@name='position'"> <xsl:copy> <xsl:apply-templates select="@*" /> <annotation> <appinfo> <domainmap domainfield="com.xis.domains.geo.GeoDomain.latLonAlt" converter="DegreesToLatLonAltConverter" parameters="lat,lon,alt" label="Position" /> </appinfo> </annotation> </xsl:copy> </xsl:when> <xsl:otherwise> <xsl:copy> <xsl:apply-templates select="@*" /> </xsl:copy> </xsl:otherwise> </xsl:choose> </xsl:for-each> </pre>

When such a transformation is used with the previous XSDL source shown above, used to type the XML document source, the result is shown in Table XIX

5 below.

Table XIX	
LEIF Data Item - tracks	
LEIF Data Item - track	
Attribute - name :	String value='FIRST'
Attribute - country :	String value='US'
Attribute - category :	String value='SUB'
Attribute - threat :	String value='HOS'
Attribute - position :	LatLonAlt <- now mapped to GeoDomain.latLonAlt
LEIF Data Item - track	
Attribute - name :	attribute value='SECOND'
Attribute - country :	String value='US'
Attribute - category :	String value='AIR'
Attribute - threat :	String value='FRI'
Attribute - position :	LatLonAlt <- now mapped to GeoDomain.latLonAlt

Note that position has been mapped to GeoDomain.latLonAlt, meaning, it is
 10 now a LatLonAlt type. The lat, lon, and alt attributes are still the same, but the
 position itself may now be dealt as a single LatLonAlt object rather than a generic

data item with three Double fields. The use of XSLT also enables the mapping of XML elements into TypeMetaData, such as the LatLonAltTypeMetaData in order have better manipulation of the data.

In one embodiment of the invention, the XIS framework has a "remote" capability enabling referral to an XIS session, or the data within it, from a remote virtual machine. Primarily, this means that data retrieval and sharing is the basis for export and collaboration. Secondly, it supports operating on data within XIS from a remote virtual machine, perhaps on information existing only at that remote virtual machine.

In another embodiment, SOAP (Simple Object Access Protocol) may be utilized. SOAP is a messaging framework that defines a protocol for the exchange of information in a decentralized, distributed environment. See, for example, the W3C Web site at <http://www.w3.org/TR/SOAP>. SOAP provides an instantiation of the remote capability, and is used to enable a general-purpose means of mating together or communicating among distributed sessions. These sessions may be two full XIS sessions on separate servers, or may comprise one XIS server and one XIS client, or may be used to interface XIS to otherwise very different software, even including software applications written in different programming languages and running on incompatible operating systems. All data and remote procedure calls are made through XML. Those skilled in the art will be understand that SOAP is an XML-based protocol. The working draft of this specification is currently being developed by the W3C.

In another embodiment, remote references are used. These references are persistable or persistent references to XIS data. In a distributed context XIS uses XML representations of the references to share, persist, or identify particular data. Embedded in the XML are custom information enabling determination of data location in files or databases, on web servers, or based on unique Ids or names within

the XIS session, or based on data which are passed to a constructor, or on an ECMA script that can re-generate the data object.

In another embodiment, the data representation in the remote distributed sessions is all based on the XIS mediation layer. All data, regardless of the initial
5 type, is passed and manipulated as sets of attributes and relationships. This greatly enhances portability and interoperability of the data. At the same time the data contains a global unique identifier, so the individual and specific source and identity of that data is maintained in spite of having been translated to a mediated format.

In another embodiment, a metadata type, e.g., "Constructor Scripts," is used.
10 This metadata type supports the generation of ECMA Script scripts to restore a TypeMetaData (metadata type) object at a later time. By executing the script, the TypeMetaData is restored to its exact state. The advantage over standard Java object serialization is control over the save/restore process, and the ability to represent the state in fairly readable text format.

15 Distributed Multi-user Collaboration

Figure 31 is a schematic diagram of the second set of facilities that allows users at remote computers to collaborate through shared views of the same data. It also shows objects and methods that may be involved in distributed collaboration facilities. Visual consumer components (e.g., INFOBEANS) are condensed and
20 compressed so that they may be efficiently accessed by a remote computer and then locally displayed. This is done using a MetaView object, which encapsulates the essential elements of a visual component to allow reconstruction on another screen with minimum requirement for data transmission. MetaViews are manipulated using the relationship handling machinery of the framework as discussed above; in
25 particular, an indirect reference to the MetaView is produced and passed around to remote consumers, who need only resolve the reference and retrieve the MetaView object if the user requests the display.

An event-handling schema is introduced for transmission of changes in the view from one computer to all others that are joined to the session. View management is accomplished via a "tree" display showing hierarchies of local views, sessions, and shared views (including properties and subsidiary data items). Views may be shared through the intuitive action of dragging and dropping them from one branch on the tree to another. Multicast protocols may be used so that MetaViews, LeifReferences (references), and Event objects may be transmitted directly from one computer to another without the need to go through a server. Alternatively, HTTP protocols may be employed if security and firewall considerations make multicasting impractical.

Metadata Server Facilities

Figure 32 shows a third set of facilities that allows remote consumer software objects 3220 to obtain additional annotations on information that is encountered along with at least a minimal self-description component, as may be found, e.g., in an XML tag name. These facilities are targeted to be useful to software agents 3202 in particular. (See Wooldridge, M., Jennings, N. (1995); "Intelligent Agents: Theory and Practice" in Knowledge Engineering Review 10:2.) For example, a Java software agent may occasionally encounter data objects or request terms, which were not anticipated when the agent was programmed. The tag name or other descriptive annotation (connector) 3206 is used as a key to request metadata and possibly domain policy methods defined within XIS from a known server, called a semantic repository 3204. The server 3204 checks its index for a match; if a match is not found, it can optionally ask a predicate logic inference engine 3208 to provide alternate names to search under (e.g., "cat, large, ferocious" is returned from the inference engine in response to a request for "lion"), which are then also compared to the index. If a match is found, the server returns the appropriate metadata and possibly domain policy method(s) to the requesting agent 3202, making use of either Sun's Jini

framework or Java's built-in provisions 3222 for network transfer and subsequent dynamic loading of classes (including, e.g., procedural components of the metadata and methods) to transfer the actual code procedures into the running agent. These transferred information can then be used by the agent 3202 to aid in handling the
5 object. For example, an agent can display an unanticipated data type to a user by calling the renderer method thus returned from the semantic repository. The semantic repository may also interface with other repositories 3212 via a data network 3210, such as a wide area network (WAN) 3210.

In the above embodiment, information metadata and even information access
10 software may be dynamically downloaded from remote sources ("semantic repositories" or "Metadata Server Facilities"), e.g., from the data source itself or from an intermediary server. This is a unique feature particularly as applied to XIS metadata, domain policies, and DSIs.

One skilled in the art will further understand that the common representation of
15 and access to information from internal APIs (mediation layer) may be exported to distributed application in support of middleware architectures for scalability or other advantages, and in support of collaboration by sharing data and data changes (permanent or transient) with other distributed collaborators.

The present invention also supports user "transfer" of objects from one context
20 (e.g., between two INFOBEANS) via drag/drop, cut/paste, and other mechanisms. Transfers of objects apply only to non-contextual information. Thus, each view has its own InfoModel that provides the contextual information for that view alone. Example 5 below further explains this idea.

Example 5

25 The Example 5 application is contained in three Java source files, HelloWorld.java, HelloWorldTranslator.java, and TestHarness.java. The source files

for this Example are explained below. The resulting outputs are shown in Figures 38A and 38B. Non-contextual information is transferred in this Example.

The Java source files in Example 5 modify those from Example 3 or 4 to demonstrate the use of drag and drop facilities within XIS, along with a more complex view. As before, changes from the previous step are marked by open and close commented braces.

The table below (Table XX) shows a code snippet that replaces a portion of code in the HelloWorld.java file shown in Example 3. The code portion below replaces the portion of code 2102 found in Figure 21 to implement Example 5.

10

Table XX	
<pre>/*{*/ private String myName; public HelloWorld(String myName) { this.myName = myName; } public String toString() { return myName; } /*}*/</pre>	

With the above modification, the HelloWorld.java file for this Example has been enhanced so that instances can be given names upon creation, and this name is returned by the toString() method. Also, the getID() method has been removed.

15

The HelloWorldTranslator.java file for this Example is similar to the HelloWorldTranslator.java file shown in Figures 22A to 22C, except that this source file has been simplified from the previous version by removing the "course" attribute.

Figures 37A to 37D list the TestHarness.java source file, where the bulk of the changes are coded. First of all, it contains an import to help lay out multiple components within a JFrame 3702, and also imports for two new XIS INFOBEANS: a Table bean and a Tree bean 3704. The tree bean displays data items and any member data items they contain, and enables users to select and perform various functions on them such as editing their properties. The table bean displays the attributes of multiple data items in a user-configurable way.

The first portion of new code 3706A and 3706B initializes several HelloWorld objects and inserts them into an array. The table and plot INFOBEANS are initialized and placed side-by-side in a JFrame. They are not loaded with any data items, but items can be dragged onto them when the application is running. The next thing is that a tree INFOBEAN is initialized and loaded with the array, and it is put up in its own JFrame.

This Example 5 application puts up two windows—the first is a Tree INFOBEAN containing several HelloWorld data items. The second has two frames, on the left a Table INFOBEAN, on the right a Chart INFOBEAN, both of which are initially empty. The table is initially completely blank because it has no information on what headings are appropriate, while the plot displays a default grid.

In the tree display, a user may select data items or open and close them, similar to the directory tree view provided by the "Windows Explorer" application. In this case, the only item that contains anything is the root, which is a special data item. Double-clicking on it opens or closes it, but it cannot itself be selected. Right clicking on an item (but not the root) brings up a menu of operations that may be performed on it. In particular, selecting 'Properties' brings up a property sheet window (as in the

previous examples) for that particular instance. A user may left-click-and-drag on data items from the property sheet window over to the table or the plot.

The table automatically sets up columns for the preferred attributes of the data items dragged onto it. The layout may be changed interactively by either dragging on the headings or their boundaries, or by right clicking on the headings. Double clicking on the attributes enables a user to edit them. Items may also be dragged from the table onto the plot as well, or in the other direction. The selection and pen color attributes are kept aligned automatically between all views of the same data item. A user exits the application by closing the tree window.

Data item commands, also known as JAF Commands (based on "JavaBeans Activation Framework") may also be implemented within XIS. The JAF API is a part of Java and provides a way to associate commands with object. By using JAF, developers may take advantage of standard Java interfaces to determine the type of an arbitrary piece of data, encapsulate access to it, discover the operations available for it, and to instantiate the appropriate bean to perform said operation(s). For example, if a browser obtained a JPEG image, JAF would enable the browser to identify that the stream of data is a JPEG image, and from that type, the browser could locate and instantiate an object that could manipulate, or view that image. JAF commands are typically used in pop-up menus or menu bars to control objects.

XIS uses JAF commands to control DSIs and perform operation on data items. Right clicking on a data item in an InfoBean pops up a menu with commands that are specific to that data item. JAF commands with XIS may implement the following:

- "Context Menus" or pop-up menus initiated by right clicking on a display that represents one (or more) data items;

- Command "menus" may be combined into a menu for multiple objects (e.g., one single command (menu item) to delete all selected objects;
- Addition of data item commands per object type by adding them to the "JAF" registry;
- Addition of file/content MIME types in the display;
- Addition to the "JAF" registry of entries that represent objects that can add any number of menu items dynamically;
- JAF commands may be added by "Menu Populators" that can add menu items for single and multiple data items;
- JAF commands may be added by the environment for other typical functions that might be performed on any object (possibly depending on capabilities/content of that data item), e.g., Copy/Paste/Cut/Delete, Hide/Show, Refers To (if the object exposed relationships to other objects which may or may not be loaded yet), Contains, Referred by, Contained by, URLs (attributes), Properties, etc.; and
- Default behavior of "Properties" command may be overridden using a "JAF entry."

Figures 38A, 38B, 38C, and 38D show how JAF is used within the framework. Figure 38A shows a menu of HelloWorld objects for selection, displayed on a screen of the computer system implementing the framework. Figure 38B shows that the HelloWorld objects are displayed through a graphing (data plot) viewer. Figure 38C shows the display of Figure 38B after the display cursor has been positioned over the bar graph, followed by a right-click. A context menu, enabling the user to select the "Properties" option, is displayed in Figure 38D. In this example, the HelloWorld

object can be manipulated by a PropertySheetInfoBean object. Thus, selecting the "Properties" option results in a PropertySheetInfoBean being instantiated showing the properties of the First HelloWorld object, as shown in Figure 38D.

Another feature provided by the XIS framework is the use of attribute aliases.

5 This enables one attribute from one domain to be substituted for an attribute in another domain. For illustration purposes, let us assume that we have two domains called Movement Domain and Physics Domain. The attribute speed in knots is defined in the Movement Domain, while the attribute velocity in meters/sec is defined in the Physics Domain. The XIS framework enables a developer to register or define
10 that the speed attribute in the Movement Domain is similar or compatible with the velocity attribute in the Physics Domain. Thus, enabling a data consumer object to ask for the Attribute from the domain that it knows about, and have the value translated from the Attribute of the DSI from the second domain.

Figure 39A, 39B, and 39C show details of an exemplary embodiment
15 explaining how an interface for the attribute alias feature described above may be defined or implemented within the XIS framework. The illustration includes class listing for objects that are used in the preferred embodiment of the framework.

Figure 40 is a block diagram of a computer that may be used to implement the framework described herein. The framework may be used in a single computer, or
20 may be used in conjunction with one or more computers that may communicate with each other over a network to share data. Those skilled in the art will appreciate that the various data objects, framework extensions, and other processes described above may be implemented with one or more computers, all of which may have a similar computer construction to that illustrated in Figure 40, or may have alternative
25 constructions consistent with the capabilities described herein.

Figure 40 shows an exemplary computer 4000 such as might comprise a computer in which an object oriented programming environment is supported to

permit the framework operations described above, and to permit the various display operations and computer processing events. Each computer 4000 operates under control of a central processor unit (CPU) 4002, such as a "Pentium" microprocessor and associated integrated circuit chips, available from Intel Corporation of Santa Clara, California, USA. A computer user can input commands and data from a keyboard and computer mouse 4004, and can view inputs and computer output at a display 4006. The display is typically a video monitor or flat panel display. The computer 4000 also includes a direct access storage device (DASD) 4008, such as a hard disk drive. The memory 4010 typically comprises volatile semiconductor random access memory (RAM). Each computer preferably includes a program product reader 4012 that accepts a program product storage device 4014, from which the program product reader can read data (and to which it can optionally write data). The program product reader can comprise, for example, a disk drive, and the program product storage device can comprise removable storage media such as a magnetic floppy disk, a CD-R disc, a CD-RW disc, or DVD disc.

The computer 4000 can communicate with other computers over a computer network 4016 (such as the Internet or an intranet) through a network interface 4018 that enables communication over a connection 4020 between the network 4016 and the computer 4000. The network interface 4018 typically comprises, for example, a Network Interface Card (NIC) or a modem that permits communications over a variety of networks.

The CPU 4002 operates under control of programming steps that are temporarily stored in the memory 4010 of the computer 4000. When the programming steps are executed, the computer performs its functions. Thus, the programming steps implement the functionality of the event tracking process described above. The programming steps can be received from the DASD 4008, through the program product storage device 4014, or through the network connection

4020. The program product storage drive 4012 can receive a program product 4014, read programming steps recorded thereon, and transfer the programming steps into the memory 4010 for execution by the CPU 4002. As noted above, the program product storage device can comprise any one of multiple removable media having recorded
5 computer-readable instructions, including magnetic floppy disks and CD-ROM storage discs. Other suitable program product storage devices can include magnetic tape and semiconductor memory chips. In this way, the processing steps necessary for operation in accordance with the invention can be embodied on a program product.

10 Alternatively, the program steps can be received into the operating memory 4010 over the network 4016. In the network method, the computer receives data including program steps into the memory 4010 through the network interface 4018 after network communication has been established over the network connection 4020 by well-known methods that will be understood by those skilled in the art without
15 further explanation. The program steps are then executed by the CPU 4002 thereby comprising a computer process.

It should be understood that the any device that supports the framework environment described above, acting as a host computer, will typically have a construction similar to that shown in Figure 40, so that details described with respect
20 to the Figure 40 computer 4000 will be understood to apply to all computers of any network system in which the framework or framework-developed applications are utilized. Alternatively, the devices can have an alternative construction, so long as the computer can communicate with the other computers and can support the functionality described herein.

25 All patents, patent applications, publications and references mentioned in the specification are indicative of the practice level of those skilled in the art to which the invention pertains. All patents, patent applications, publications and references are

herein incorporated by reference to the same extent as if each individual patent, patent application, publication or reference was specifically and individually indicated to be incorporated by reference.

One skilled in the art should readily appreciate that the present invention is
5 well adapted to carry out the objects and obtain the ends and advantages mentioned,
as well as those inherent therein. The specific embodiments described herein as
presently representative of preferred embodiments are exemplary and are not intended
as limitations on the scope of the invention. Changes therein and other uses will
occur to those skilled in the art which are encompassed within the spirit of the
10 invention are defined by the scope of the claims. It will be readily apparent to one
skilled in the art that modifications may be made to the invention disclosed herein
without departing from the scope and spirit of the invention.

The present invention has been described above in terms of presently preferred
embodiments so that an understanding of the present invention can be conveyed.
15 There are, however, many configurations for information system frameworks not
specifically described herein but with which the present invention is applicable. The
present invention should therefore not be seen as limited to the particular
embodiments described herein, but rather, it should be understood that the present
invention has wide applicability with respect to information systems generally. All
20 modifications, variations, or equivalent arrangements and implementations that are
within the scope of the attached claims should therefore be considered within the
scope of the invention.

CLAIMS

We claim:

5

1. An extensible framework for use in a computer system that supports an object oriented programming environment and includes a memory in which data objects can be stored, the framework comprising a set of object classes that can be extended using object oriented principles to define an information handling
10 application with data objects comprising a class of data source objects and a class of data consumer objects, and a mediation layer that defines an interface between the data source objects and the data consumer objects to permit data communications between the two data object classes, such that the class configuration of the data source objects can be specified independently of the class configuration of the data
15 consumer objects.

2. A computer system that supports an object oriented programming environment and includes access to data storage containing specifications for a set of object classes that can be extended using object oriented principles to define an
20 information handling application with data objects comprising a class of data source objects and a class of data consumer objects, and a mediation layer that defines an interface between the data source objects and the data consumer objects to permit data information exchange between the two data object classes, such that the class configuration of the data source objects can be specified independently of the class
25 configuration of the data consumer objects.

3. A computer system as defined in claim 2, wherein the mediation layer defines a data interface that provides an information exchange standard between the data source objects and the data consumer objects.

5 4. A computer system as defined in claim 3, wherein the data interface includes at least one attribute/metadata object class that specifies attributes, including metadata, and that provide declarative and procedural information relating to attributes in said data object.

10 5. A computer system as defined in claim 3, wherein the data interface includes at least one Relationship object class that specifies relationships between data source class objects.

15 6. A computer system as defined in claim 3, wherein the data interface includes at least one Domain Policy object class that specifies a group of related attributes, methods, and semantic information that indicates data processing to be available for the specified attributes, including the specific intent of said attributes.

20 7. A computer system as defined in claim 3, wherein the data interface includes at least one Event Change object class that specifies change event registration for detecting changes in the data objects.

25 8. A computer system as defined in claim 4, further including a TypeMetaData class of objects that specify semantic information indicating data processing to be available for specified attributes and intended use of the attributes of a data object.

9. A computer system as defined in claim 8, further including a TypeIO class of objects that define a desired format of a data object attribute specified by a TypeMetaData class.

5 10. A computer system as defined in claim 6, wherein an attribute alias may be specified by a user to indicate any data source attributes of an attribute domain that may be substituted with attributes of a different attribute domain.

11. A computer system as defined in claim 4, wherein said
10 attribute/metadata interface provides input/output functions for display and editing of the attribute during runtime.

12. A computer system as defined in claim 4, further including a
15 FieldMetaData attribute that specifies processing criteria for an attribute of a data object.

13. A computer system as defined in claim 4, wherein the processing criteria relates to a selected attribute of the data object for display processing.

20 14. A computer system as defined in claim 5, wherein the indicated relationship may be resolved using a version of the Relationship object in cache of the computer system.

15 15. A computer system as defined in claim 2, further including a collaboration facility that permits a user at a network computer remote from the computer system to share a view of a data object at the remote network computer.

16. A computer system as defined in claim 15, wherein the collaboration facility comprises a MetaView class of data objects that lightweight encapsulate elements of a data object for visual reconstruction on a display screen of the remote network computer.

5

17. A computer system as defined in claim 4, wherein the data interface provides an override function that can override default metadata of the data object attributes.

10

18. A computer system as defined in claim 6, wherein the data interface provides an override function that can override default metadata of the data object attributes.

15

19. A computer system as defined in claim 3, wherein data objects provided through the data interface comprise information that specifies data attributes, relationships, event broadcasting of changes in all registered data consumer objects, metadata for attributes that provide declarative and procedural information and/or input/output functions for display and editing, and related data items, and wherein the data objects can store references without direct access to a related data item.

20

20. A computer system as defined in claim 19, wherein the attributes and relationships include contextual metadata, and the events and methods are defined with respect to the contextual metadata.

25

21. A computer system as defined in claim 3, wherein the mediation layer further includes object classes that include one or more methods that provide data exposure, and further including a Data Source Interface object class of data objects

with methods that automatically determine which data exposure method will be used for data communications between the data source objects and data consumer objects.

22. A computer system as defined in claim 21, wherein the determined data
5 exposure method comprises data object reflection and introspection.

23. A computer system as defined in claim 21, wherein the data exposure
methods include a data source object method that provides a standard interface, a data
source object and a Translator object that maps an alternate data interface to the
10 standard interface, and a data source object that is inspected by the computer system
to determine available data fields which are thereby exposed to data consumer objects
by the standard interface.

24. A computer system as defined in claim 3, wherein the data interface of
15 the mediation layer provides a wrapper for the data objects.

25. A computer system as defined in claim 24, wherein data exchange
occurs without providing contextual data characteristics to a receiving data object.

20 26. A computer system as defined in claim 25, wherein the contextual data
characteristics are supplied by the data object wrapper.

27. A computer system as defined in claim 3, wherein data exchange occurs
without providing contextual data characteristics to a receiving data object.

25 28. A computer system as defined in claim 2, further including a plug-in
manager object class that integrates service components into the application.

29. A computer system as defined in claim 28, wherein data service components interact with other components of the application to determine if a service is required and to determine parameters that are to be interchanged.

5

30. A computer system as defined in claim 29, wherein the service determination occurs upon introduction of newly loaded components.

31. A computer system as defined in claim 3, further including a translator
10 facility that translates information concerning a data object into (1) attributes and metadata that define groups of data source and data consumer class attributes and (2) an object class that specifies at least one domain method for defining groups of data object attributes.

15 32. A computer system as defined in claim 24, further including an InfoModel data object class.

33. A computer system as defined in claim 27, wherein said contextual data characteristics are supplied by an InfoModel object.

20

34. A computer system as defined in claim 33, wherein the InfoModel object receives a data source object for exposure to the data consumer objects, selects available attributes and methods of the received data source object for data characterization, and determines if the addition of attributes or methods to the data
25 source object are appropriate.

35. A computer system as defined in claim 6, wherein the domain is defined using XML schema.

36. A method of communicating data in a computer system that supports an object oriented programming environment and includes data storage or access to data storage containing data objects and specifications for a set of object classes that can be extended using object oriented principles to define an information handling application, the method comprising:

receiving one or more data objects comprising a class of data source objects;
10 representing the data source objects in accordance with an Information Model class of objects of a mediation layer that defines a data interface between the data source objects and a class of data consumer objects;

wherein the Information Model objects include methods that permit data communications or information exchange between the two data object classes, such
15 that the class configuration of the data source objects can be specified independently of the class configuration of the data consumer objects.

37. A method as defined in claim 36, further including specifying at least one attribute/metadata object class that specifies attributes, including metadata, and
20 that provides declarative and procedural information relating to attributes in said data object.

38. A method as defined in claim 36, further including specifying a Relationship object class or classes that specify relationships between data source
25 objects.

39. A method as defined in claim 36, further including specifying a Domain Policy or Domain object class or classes that specify a group of related attributes, methods, and semantic information that indicates data processing to be available for the specified attributes, including the specific intent of said attributes.

5

40. A method as defined in claim 36, further including specifying an Event Change object class or classes that specify change event registration for detecting changes in the data objects.

10

41. A method as defined in claim 37, further including specifying a TypeMetaData class of objects that specify semantic information indicating data processing to be available for specified attributes and intended use of the attributes of a data object.

15

42. A method as defined in claim 41, further including specifying a TypeIO class of objects that define a desired format of a data object attribute specified by a TypeMetaData class.

20

43. A method as defined in claim 39, further including specifying an attribute alias by a user to indicate any data source attributes of an attribute domain that may be substituted with attributes of a different attribute domain.

25

44. A method as defined in claim 37, wherein the attribute/metadata interface provides input/output functions for display and editing of the attribute during runtime.

45. A method as defined in claim 37, wherein the attribute/metadata object class is a FieldMetaData attribute that specifies processing criteria for an attribute of a data object.

5 46. A method as defined in claim 37, wherein the processing criteria relates to a selected attribute of the data object for display processing.

47. A method as defined in claim 38, wherein the indicated relationship may be resolved using a version of the Relationship object in cache of the computer
10 system.

48. A method as defined in claim 36, further including providing a collaboration facility that permits a user at a network computer remote from the computer system to share a view of a data object at the remote network computer.

15 49. A method as defined in claim 48, wherein the collaboration facility comprises a MetaView class of data objects that lightweight encapsulate elements of a data object for visual reconstruction on a display screen of the remote network computer.

20 50. A method as defined in claim 37, wherein the data interface provides an override function that can override default metadata of the data object attributes.

51. A method as defined in claim 39, wherein the data interface provides an
25 override function that can override default metadata of the data object attributes.

52. A method as defined in claim 36, wherein data objects provided through the data interface comprise information that specifies data attributes, relationships, event broadcasting of changes in all registered data consumer objects, metadata for attributes that provide declarative and procedural information and/or input/output functions for display and editing, and related data items, and wherein the data objects can store references without direct access to a related data item.

53. A method as defined in claim 52, wherein the attributes and relationships include contextual metadata, and the events and methods are defined with respect to the contextual metadata.

54. A method as defined in claim 36, wherein the mediation layer further includes object classes that include one or more methods that provide data exposure, and further including a Data Source Interface object class of data objects with methods that automatically determine which data exposure method will be used for data communications between the data source objects and data consumer objects.

55. A method as defined in claim 54, wherein the determined data exposure method comprises data object reflection and introspection.

56. A method as defined in claim 54, wherein the data exposure methods include a data source object method that provides a standard interface, a data source object and a Translator object that maps an alternate data interface to the standard interface, and a data source object that is inspected by the computer system to determine available data fields which are thereby exposed to data consumer objects by the standard interface.

57. A method as defined in claim 36, wherein the data interface of the mediation layer provides a wrapper for the data objects.

58. A method as defined in claim 57, wherein data exchange occurs without
5 providing contextual data characteristics to a receiving data object.

59. A method as defined in claim 58, wherein the contextual data characteristics are supplied by the data object wrapper.

60. A method as defined in claim 36, wherein data exchange occurs without
10 providing contextual data characteristics to a receiving data object.

61. A method as defined in claim 36, further including a plug-in manager object class that integrates service components into the application.
15

62. A method as defined in claim 61, wherein data service components interact with other components of the application to determine if a service is required and to determine parameters that are to be interchanged.

63. A method as defined in claim 62, wherein the service determination
20 occurs upon introduction of newly loaded components.

64. A method as defined in claim 36, further including a translator facility that translates information concerning a data object into (1) attributes and metadata
25 that define groups of data source and data consumer class attributes and (2) an object class that specifies at least one domain method for defining groups of data object attributes.

65. A method as defined in claim 57, further including an InfoModel data object class.

5 66. A method as defined in claim 60, wherein said contextual data characteristics are supplied by an InfoModel object.

67. A method as defined in claim 66, wherein the InfoModel object receives a data source object for exposure to the data consumer objects, selects available
10 attributes and methods of the received data source object for data characterization, and determines if the addition of attributes or methods to the data source object are appropriate.

68. A method as defined in claim 39, wherein the domain is defined using
15 XML schema.

69. A computer system that supports an object oriented programming environment, the system comprising:

a processor that executes program instructions to provide the object oriented
20 programming environment; and

data store means for providing program instructions containing specifications for a set of object classes that can be extended using object oriented principles to define an information handling application with data objects comprising a class of data source objects and a class of data consumer objects, and a mediation layer that
25 defines an interface between the data source objects and the data consumer objects to permit data information exchange between the two data object classes, such that the

class configuration of the data source objects can be specified independently of the class configuration of the data consumer objects.

70. A computer system as defined in claim 69, wherein the mediation layer
5 defines a data interface that provides an information exchange standard between the data source objects and the data consumer objects.

71. A computer system as defined in claim 70, wherein the data interface
includes at least one attribute/metadata object class that specifies attributes, including
10 metadata, and that provide declarative and procedural information relating to attributes in said data object.

72. A computer system as defined in claim 70, wherein the data interface
includes at least one Relationship object class that specifies relationships between
15 data source class objects.

73. A computer system as defined in claim 70, wherein the data interface
includes at least one Domain Policy object class that specifies a group of related
attributes, methods, and semantic information that indicates data processing to be
20 available for the specified attributes, including the specific intent of said attributes.

74. A computer system as defined in claim 70, wherein the data interface
includes at least one Event Change object class that specifies change event
registration for detecting changes in the data objects.

25

75. A program product for use in a computer system that executes program
instructions recorded in a computer-readable media to perform a method for

information exchange in a computer system that supports an object oriented programming environment and includes access to data storage containing data objects, the program product comprising:

a recordable media; and

- 5 a program product of computer-readable instructions executable by the computer system to perform a method comprising:

receiving data specifications for a set of object classes that can be extended using object oriented principles to define an information handling application, wherein the extended objects provide an information handling application that can
10 receive one or more data objects comprising a class of data source objects, and represent the data source objects in accordance with an Information Model class of objects of a mediation layer that defines a data interface between the data source objects and a class of data consumer objects;

- wherein the Information Model objects include methods that permit data
15 communications or information exchange between the two data object classes, such that the class configuration of the data source objects can be specified independently of the class configuration of the data consumer objects.

76. A program product as defined in claim 75, wherein the method further
20 includes specifying at least one attribute/metadata object class that specifies attributes, including metadata, and that provides declarative and procedural information relating to attributes in said data object.

77. A program product as defined in claim 75, wherein the method further
25 includes specifying a Relationship object class or classes that specify relationships between data source objects.

78. A program product as defined in claim 75, wherein the method further includes specifying a Domain Policy or Domain object class or classes that specify a group of related attributes, methods, and semantic information that indicates data processing to be available for the specified attributes, including the specific intent of said attributes.

79. A program product as defined in claim 75, further including specifying an Event Change object class or classes that specify change event registration for detecting changes in the data objects.

10

80. A method of operating a computer system to execute an information handling application in an object oriented programming environment and to receive data from a class of data source objects and provide a class of data consumer objects with data, the method comprising:

15 selecting available attributes and methods of a data source object for characterization;

determining attributes or methods to be added to the data source object during application execution; and

20 receiving the data source object and determining whether to expose or hide available attributes and methods of the object from data consumer objects, wherein data communications between the two data object classes is supported, such that the class configuration of the data source objects can be specified independently of the class configuration of the data consumer objects.

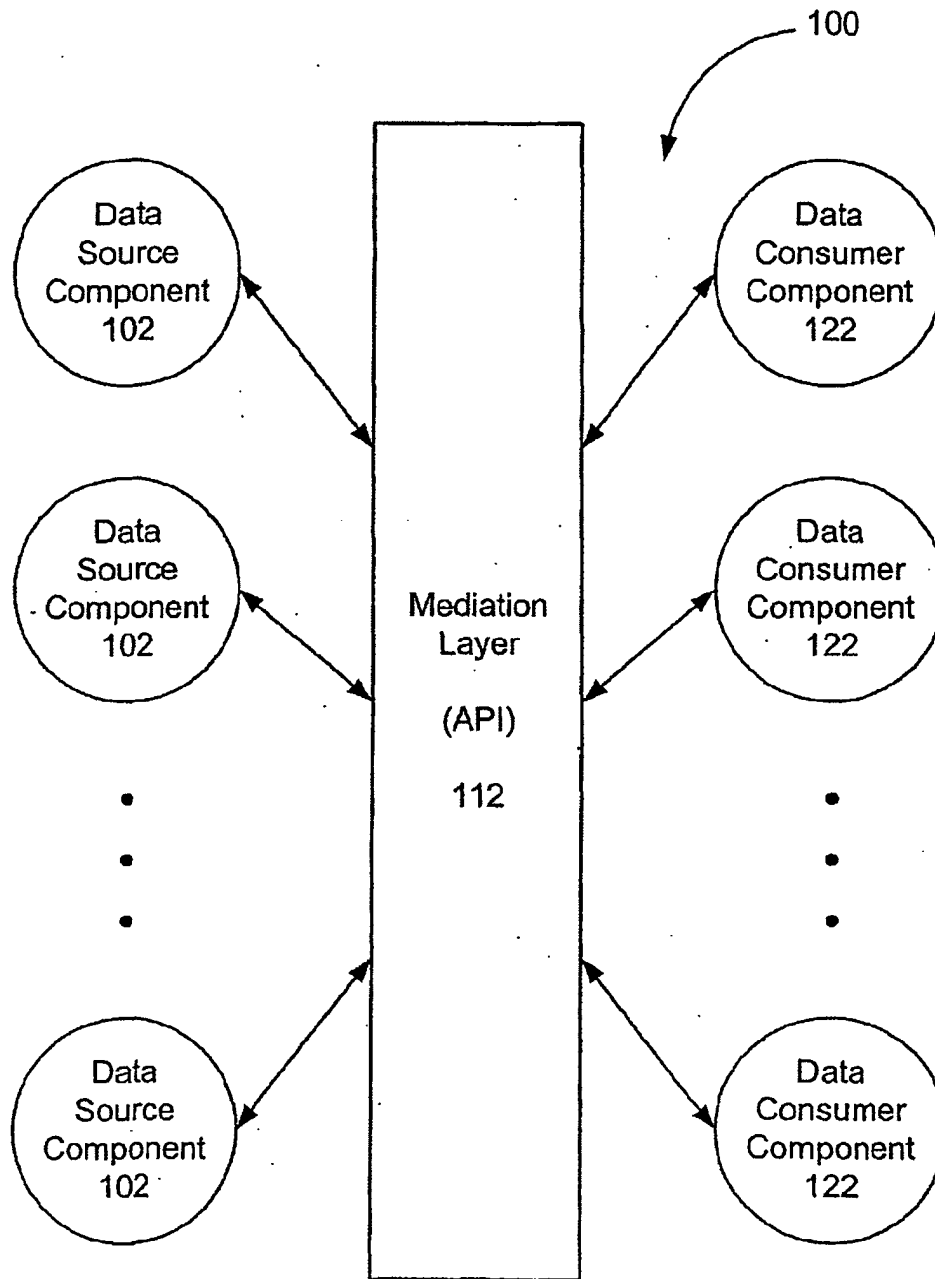


Figure 1

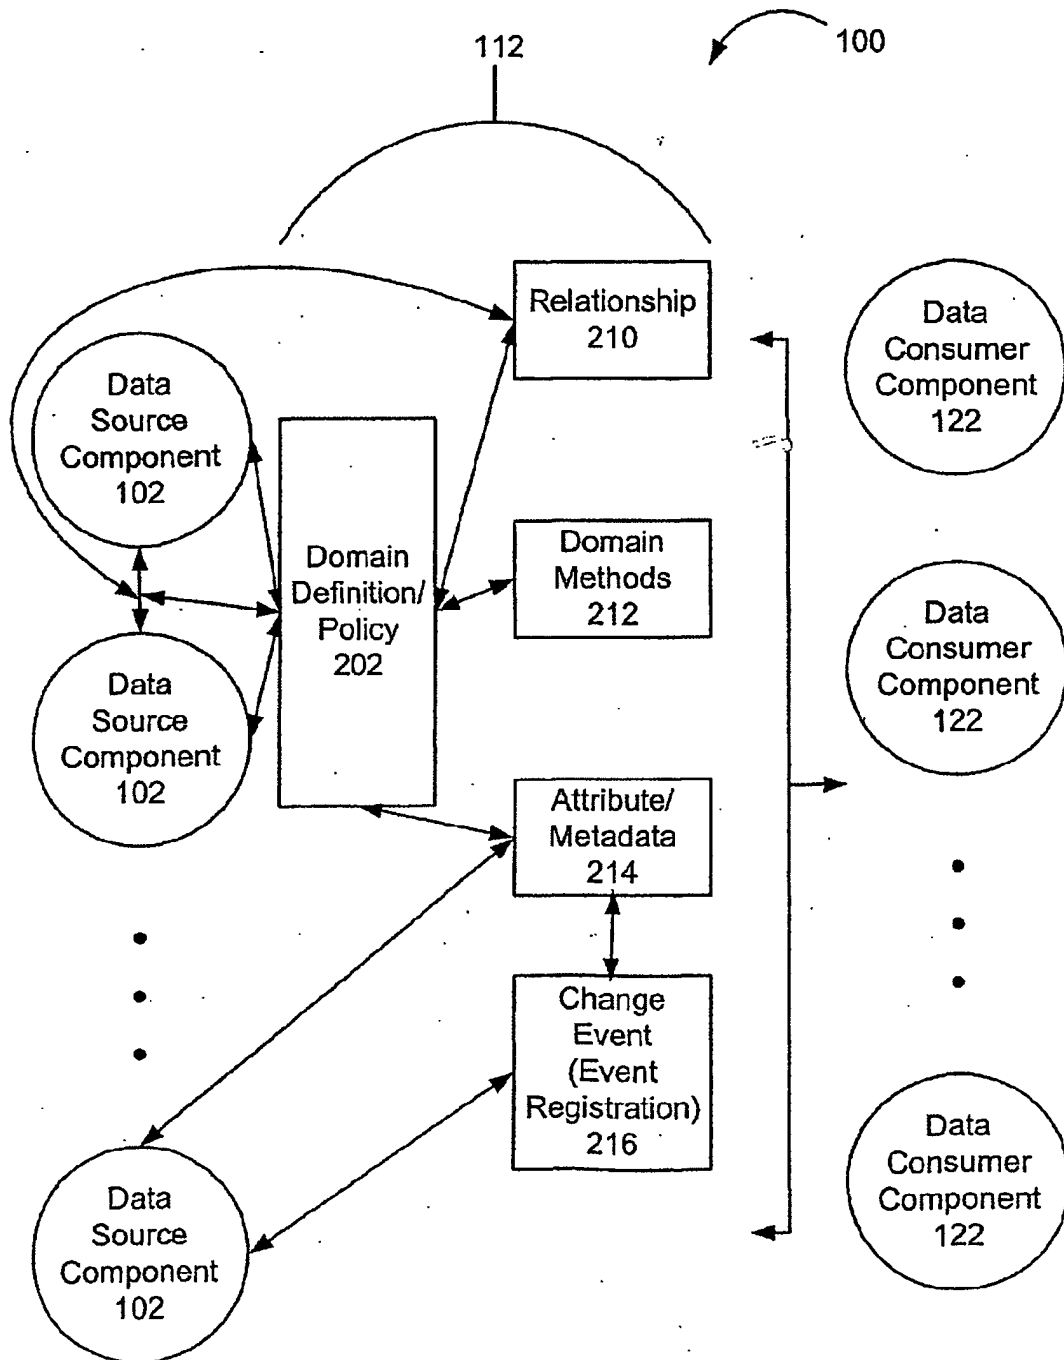


Figure 2

3/90

Figure 3A


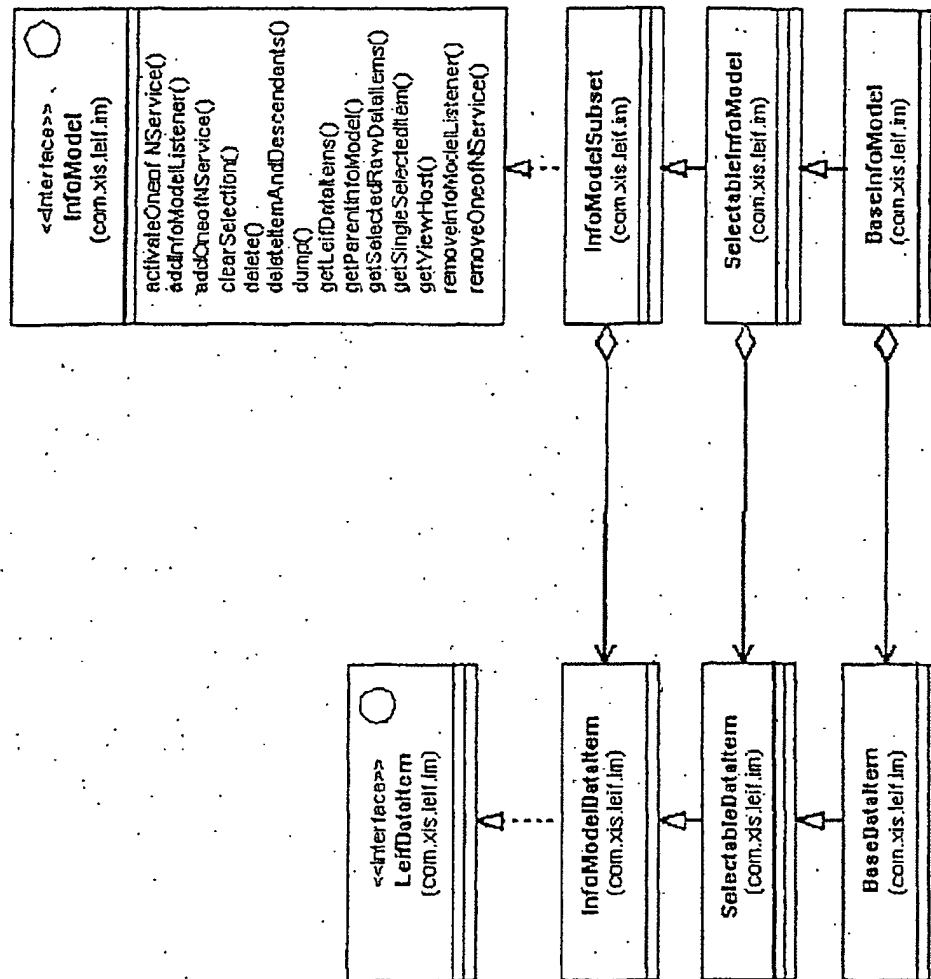


FIGURE 3b



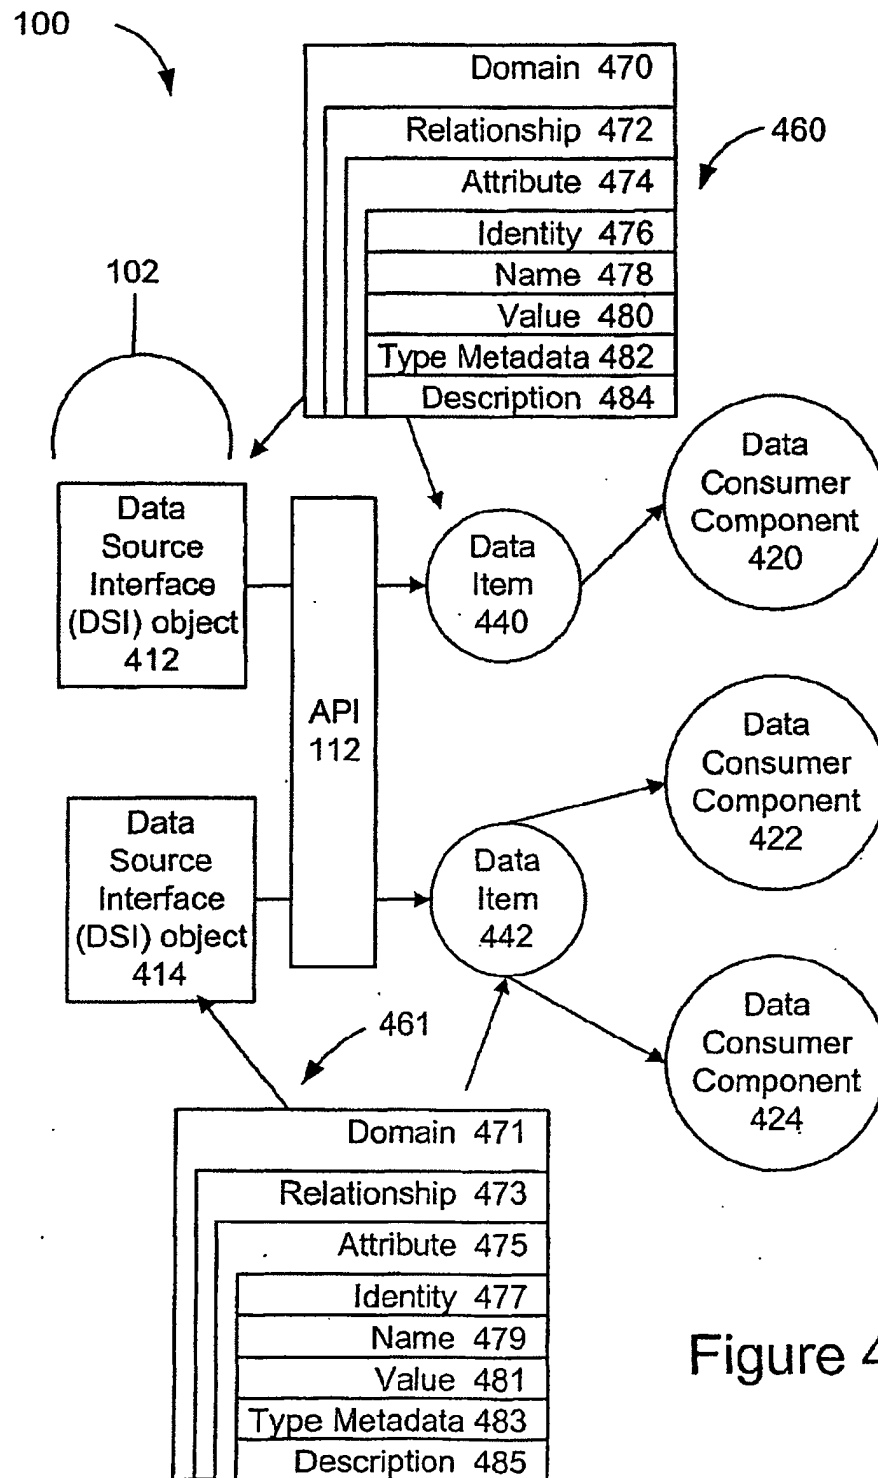


Figure 4

Figure 5

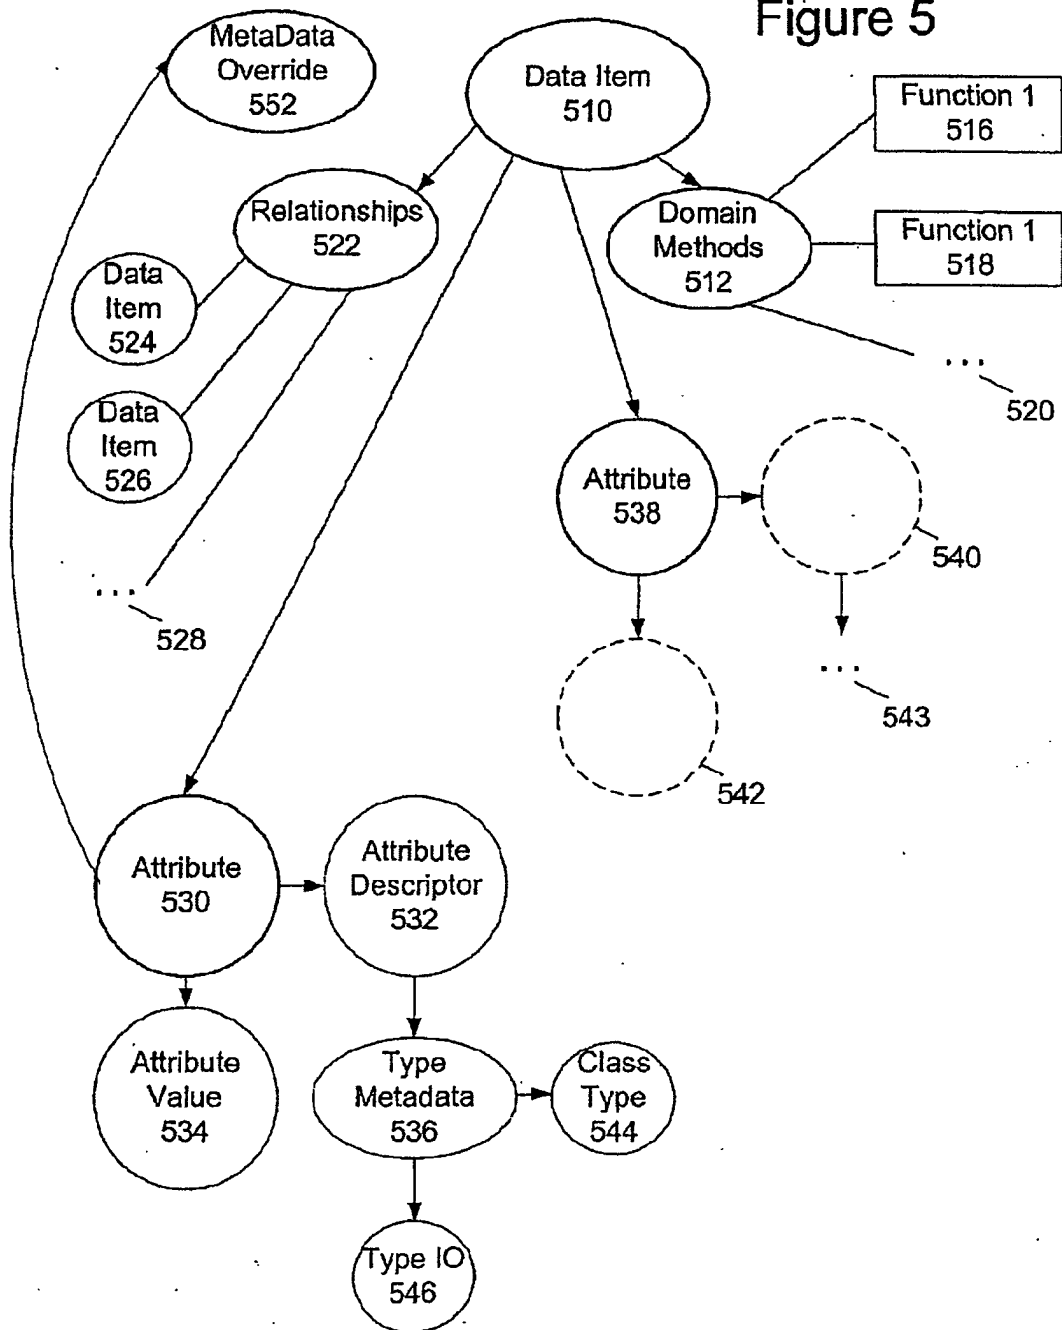


Figure 6A

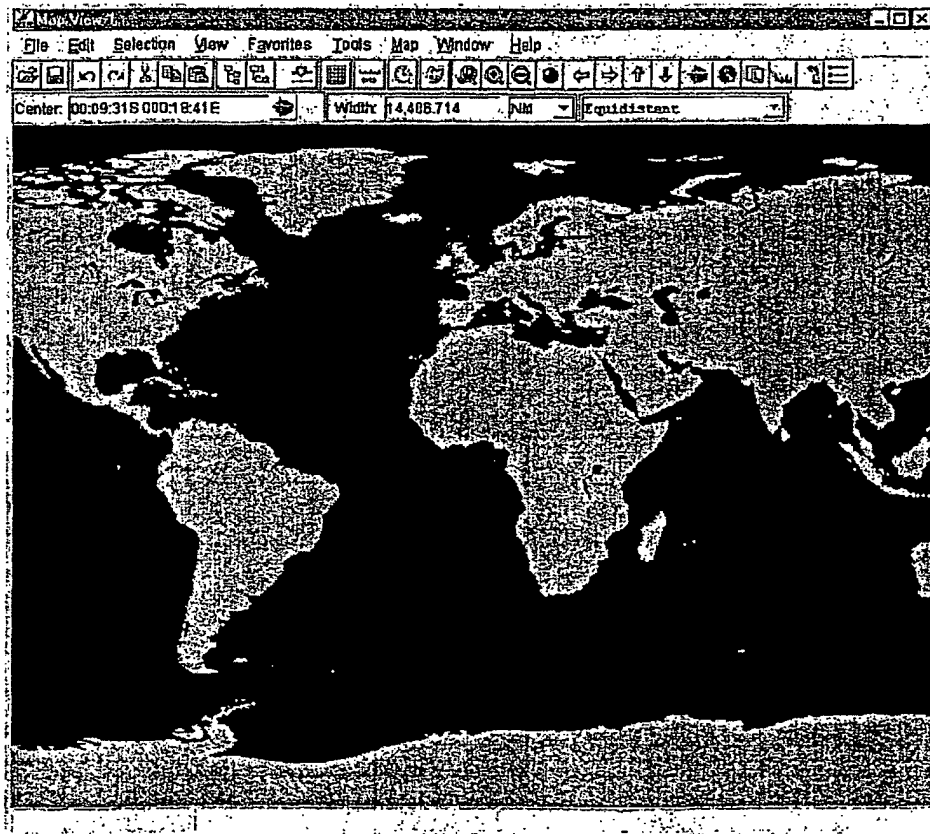


Figure 6B




Figure 6C

People Source - XIS Explorer

File Edit Selection View Favorites Tools Explorer Window Help

Explorer Contents

- Orders
- Computer
- People Source

Name	Address	City	State	Zip Code	Tr
Susan Sara...	85 Elm Street	Stratford	IL	23935	(899)4
Tim Curry	2415 Camino Del Mar	San Diego	CA	92381	(858)E
Bany Bost...	122 Thames Street	Newport	RI	14285	(401)E
Richard O...	33 Lockwood Avenue	Preston	PA	15003	(239)E
Patricia Quinn	1288 Huntington Avenue	Boston	MA	02119	(817)E
Nell Campbell	444 Mix Avenue	Hamden	CT	03419	(203)E
Jonathan ...	1 Quincy Place	Quincy	MA	02088	(817)E
Peter Hin...	238 Milsure Street	Sedona	AZ	84921	(851)E
Meat Loaf	13 Gwendal Avenue	Los Angeles	CA	92114	(818)E
Charles Gray	3928 General Street	Litleton	NH	03561	(603)E
Jeremy N...	118 Cornwall Place	Stallion	TX	43208	(838)E
Hilary Labow	9 Overlook Drive	Medway	MA	02053	(508)E
Sandra Rix	44 End Lane	Templeton	WA	84937	(828)E

Preferred Attributes

Property	Value
-Name	People Source
-Children Are Same Type	<input type="checkbox"/>
-Select Expansion	
-Name	Susan Sarandon
-Address	85 Elm Street
-City	Stratford
-State	IL
-Zip Code	23935

Apply Reset

Figure 6D

Order #70 - XIS Explorer

File Edit Selection View Favorites Tools Explorer Window Help

Orders

- Order #67
- Order #70
- Order #72
- Order #74
- Order #76
- Order #83
- Order #87
- Order #84
- Order #88
- Order #105
- Order #109
- Order #117
- Order #118
- Order #122
- Order #128
- Order #132
- Order #141
- Order #143
- Order #145
- Order #155
- Order #159
- Order #168
- Order #170
- Order #173
- Order #177
- Order #180
- Order #200
- Order #204
- Order #205
- Order #213
- Order #214
- Order #218
- Order #220
- Order #223

Name	Quantity	Retail Price (USD)	Product ID
Basketball (#1)	7	\$4.95	1
Soccer ball (#3)	12	\$12.95	3
Tether ball (#9)	14	\$9.65	9

Preferred Attributes

Property	Value
Name	Order #70
Number of Items	33
Total (USD)	\$27.55
Payment Method	Check
Processing Time (DAY)	3
Location	42:27:58N 083:11:02W
Pen Color	<input type="checkbox"/>
Zipcode	48237

Apply Reset

11/90

Figure 6E




Figure 6F




Figure 7

The screenshot shows a dialog box titled "Property Sheet - Commander". It contains a table with the following data:

Property	Value
SSN	555-55-5555
Display Name	Commander
DOB	Sep 4, 2010 5:00:02 AM
Name	Commander

At the bottom of the dialog box are four buttons: "Ok", "Cancel", "Reset", and "Apply".

Figure 8

```
Person
(com.XIS.test)

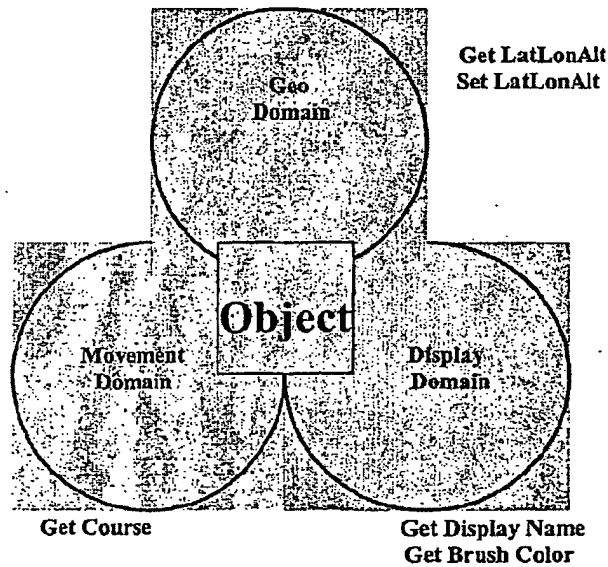
-dob : Date
-name : String
-ssn : String

+getDOB() : Date
+getName() : String
+getSSN() : String
+Person(name : String)
+Person()
+setDOB(dtg : Date) : void
+setName(newName : String) : void
+setSSN(newSSN : String) : void
+toString() : String
```

Figure 9

getAttributes()	References	Referrers	Members
com.xls.domains.display.DisplayDomain.displayName	None	None	None
com.xls.test.Person.DOB			
com.xls.test.Person.SSN			
com.xls.test.Person.name			

Figure 10



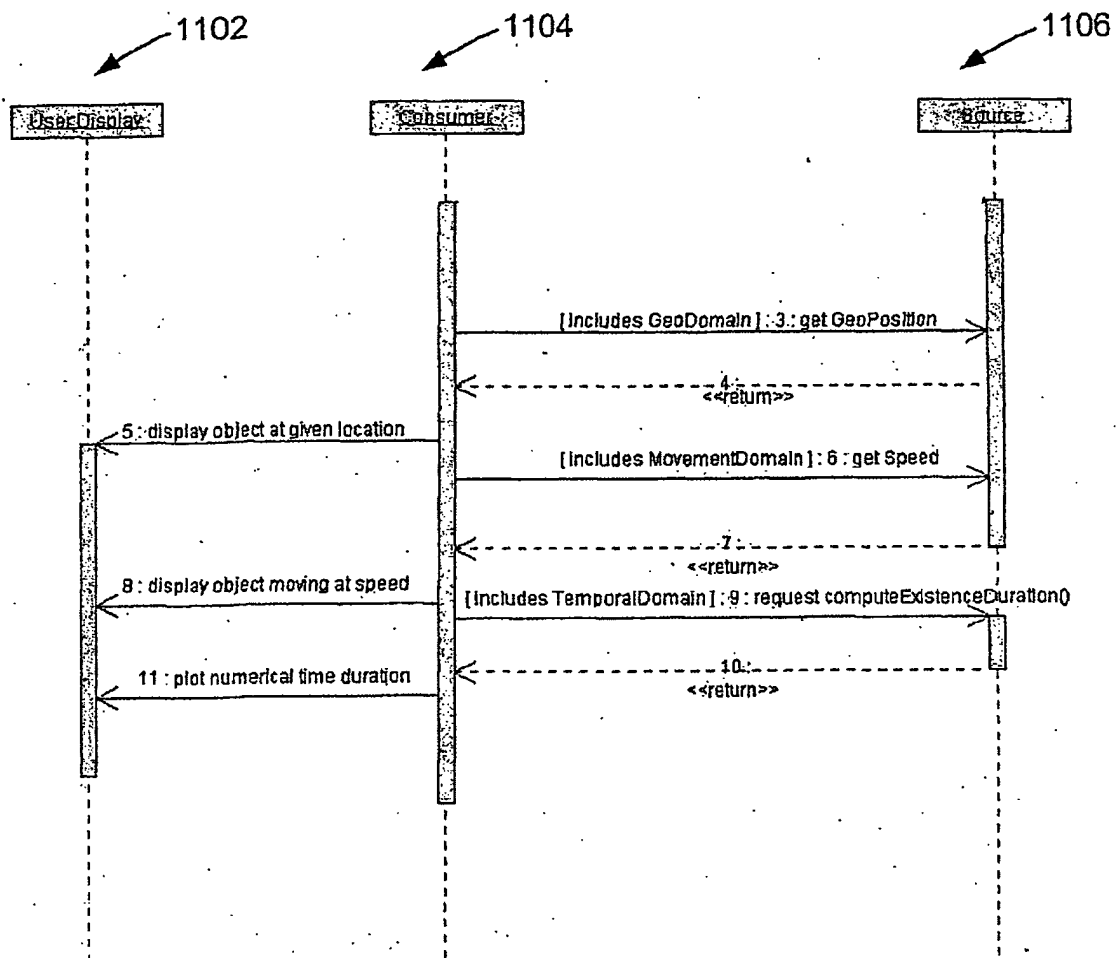


Figure 11

Package com.xis.types

This package contains classes that provide several standard TypeMetaData classes for describing types and their constraints, and for rendering and editing values of those types.

See:

Description

Interface Summary	
<u>DataTest</u>	The DataTest interface specifies methods for Object validation.
<u>HTMLTypeIO</u>	This interface defines the IO for HTML.
<u>SummaryFunction</u>	The SummaryFunction interface defines generic summary functionality based upon provided input data values.
<u>SwingTypeIO</u>	The SwingTypeIO interface allows for the use of both swing editors, allowing swing components to edit an object, and swing renderers, which know how to render these objects in a swing environment.
<u>TextTypeIO</u>	The TextTypeIO interface provides a means of formatting objects in a textual fashion, as well as parsing text from which an object is created.
<u>TypedValue</u>	The TypedValue interface is used to hold an object that carries its own TypeMetaData with it.
<u>TypeEditor</u>	The TypeEditor interface defines methods for editing attributes provided by the types implemented within this package.
<u>TypeIO</u>	The TypeIO interface provides a common base from which other TypeIOs can extend, such as HTMLTypeIO, SwingTypeIO, etc.
<u>TypeMetaData</u>	The TypeMetaData interface defines generic type accessors for object comparing, editing, formatting, rendering, and validation.
<u>TypeMetaDataFactory</u>	The TypeMetaDataFactory interface defines a class that can create TypeMetaData for a given Class Type.
<u>TypeRenderer</u>	The TypeRenderer interface defines methods for rendering attributes provided by the types implemented within this package.
<u>ValidTestProxy</u>	The ValidTestProxy interface
<u>WMLTypeEditor</u>	The WMLTypeEditor interface defines methods for rendering attributes provided by the types implemented within this package.
<u>WMLTypeIO</u>	This interface defines the IO for WML format.
<u>WMLTypeRenderer</u>	The WMLTypeRenderer interface defines methods for rendering attributes provided by the types implemented within this package.
<u>XMLTypeIO</u>	The XMLTypeIO interface provides a means of formatting objects in a XML textual fashion, as well as parsing XML text for creating an object.

FIG. 12A

Class Summary	
<u>AbstractDataTest</u>	The AbstractDataTest class provides a default implementation of <i>test (Object)</i>
<u>AbstractTypeMetaData</u>	AbstractTypeMetaData provides a partial implementation of TypeMetaData to relieve the XIS developer from explicitly implementing irrelevant methods.
<u>AreaOfUncertaintyTypeMetaData</u>	AreaOfUncertaintyTypeMetaData is a type object that supports <i>AreaOfUncertainty</i> objects.
<u>ArrayListTypeMetaData</u>	ArrayListTypeMetaData is a type object that supports <i>java.util.ArrayList</i> objects.
<u>ArrayTypeMetaData</u>	ArrayTypeMetaData is a type object that supports <i>java.lang.reflect.Array</i> objects.
<u>BeanTypeMetaData</u>	BeanTypeMetaData is a type object that supports <i>java beans</i> objects.
<u>BooleanTypeMetaData</u>	The BooleanTypeMetaData is a type object that supports <i>Boolean</i> objects.
<u>BooleanTypeMetaDataFactory</u>	A BooleanTypeMetaDataFactory can create a BooleanTypeMetaData for given booleans.
<u>CachedTypeMetaData</u>	CachedTypeMetaData is a type object that simply delegates all TypeMetaData calls to another TypeMetaData.
<u>ClassificationTypeMetaData</u>	ClassificationTypeMetaData is a type object that supports <i>Classification</i> objects.
<u>CollectionsTypeMetaData</u>	CollectionsTypeMetaData is a type object that supports <i>Collection</i> objects.
<u>ColorTypeMetaData</u>	The ColorTypeMetaData class implements TypeMetaData for Color objects.
<u>ColorTypeMetaDataFactory</u>	A ColorTypeMetaDataFactory can create a ColorTypeMetaData for a given Color object.
<u>ConversionNumericTypeMetaData</u>	A generic NumericTypeMetaData for converting from one unit to another.
<u>CurrencyTypeMetaData</u>	CurrencyTypeMetaData is a type object that supports <i>Number</i> objects that represent Currency values.
<u>DateTimeTypeMetaData</u>	DateTimeTypeMetaData is a type object that supports <i>Date</i> objects.
<u>DateTimeTypeMetaDataFactory</u>	A DateTimeTypeMetaDataFactory can create a DateTimeTypeMetaData for given Date objects.
<u>DiscreteRangeStringTypeMetaData</u>	DiscreteRangeStringTypeMetaData is a type object that supports <i>String</i> objects with discrete ranges.
<u>DisplayLabelTypeMetaData</u>	DisplayLabelTypeMetaData is a type object that supports <i>DisplayLabel</i> objects.
<u>DTGTypeMetaData</u>	DTGTypeMetaData is a type object that supports <i>Date</i> objects.
<u>EnumerationType</u>	Class to implement an enumeration in Java.
<u>EnumerationTypeMetaData</u>	The EnumerationTypeMetaData class is used to represent integer constants as strings to the user.
<u>FontTypeMetaData</u>	FontTypeMetaData is a type object that supports <i>Font</i> objects.
<u>HashMapTypeMetaData</u>	HashMapTypeMetaData is a type object that supports <i>java.util.HashMap</i> objects.

FIG. 12 B

<u>HashSetTypeMetaData</u>	HashSetTypeMetaData is a type object that supports <i>java.util.HashSet</i> objects.
<u>IconShapeTypeMetaData</u>	IconShapeTypeMetaData is a type object that supports <i>IconShape</i> objects.
<u>IconTypeMetaData</u>	IconTypeMetaData is a type object that supports <i>Icon</i> objects.
<u>LinkedListTypeMetaData</u>	LinkedListTypeMetaData is a type object that supports <i>java.util.LinkedList</i> objects.
<u>ListTypeMetaData</u>	ListTypeMetaData is a type object that supports <i>java.util.List</i> objects.
<u>MouseMapProxy</u>	The MouseMapProxy class allows TypeEditors access to the map without requiring them to having any compile time knowledge of the map's existence.
<u>NumberComparator</u>	An implementation of the Comparator interface that compares two objects that extend Number, or that both implement Comparable, with the class of one assignable from the other.
<u>NumericTypeMetaData</u>	NumericTypeMetaData is a type object that supports <i>Number</i> objects.
<u>NumericTypeMetaDataFactory</u>	A NumericTypeMetaDataFactory can create NumericTypeMetaData for a given class type or method return type.
<u>ObjectTypeMetaData</u>	ObjectTypeMetaData is a type that supports a simple <i>Object</i> and is provided to quickly add arbitrary attribute types to a Data Source Interface without writing a more specific type handler.
<u>PercentTypeMetaData</u>	PercentTypeMetaData is a type object that supports <i>Number</i> objects that represent Percent(%) values.
<u>ProbabilityTypeMetaData</u>	ProbabilityTypeMetaData is a type object that supports <i>Number</i> objects that represent Probability values.
<u>RenamedTypeMetaData</u>	RenamedTypeMetaData simply delegates all TypeMetaData calls to another TypeMetaData except for the <i>getName()</i> , which is overridden with the given value.
<u>ResizedTextTypeMetaData</u>	ResizedTextTypeMetaData simply delegates all TypeMetaData calls to another TypeMetaData except for the <i>getPixelWidth()</i> , which is overridden with the given value.
<u>Resources</u>	The Resources class is automatically generated and must be public, but it is intended to be used only by Java's internationalization support classes.
<u>SmartDurationTypeMetaData</u>	A SmartDurationTypeMetaData for converting from one time unit to another based on the magnitude of the duration value.
<u>StringTypeMetaData</u>	StringTypeMetaData is a type object that supports <i>String</i> objects.
<u>StringTypeMetaDataFactory</u>	A StringTypeMetaDataFactory can create a StringTypeMetaData for a given String object.
<u>TextTypeMetaData</u>	TextTypeMetaData is a type object that supports <i>Text</i> objects.
<u>TypedValueTypeMetaData</u>	TypedValueTypeMetaData is a type object that supports supports <i>TypedValue</i> objects.
<u>TypeEditorBeanContextChildSupport</u>	The TypeEditorBeanContextChildSupport class handles most of the responsibilities of a TypeEditor and a BeanContextChild.
<u>TypeEditorSupport</u>	The TypeEditorSupport class handles most of the responsibilities of a TypeEditor.
<u>TypeIOPluggableService</u>	The TypeIOPluggableService class is responsible for loading TypeIOs.
<u>TypeIORegistry</u>	The TypeIORegistry class is a registry for different implementations of TypeIO classes.

FIG. 12C

<u>TypeIOSupport</u>	The TypeIOSupport provides support for TypeMetadata's getTypeIO methods.
<u>TypeMetadataDelegator</u>	TypeMetadataDelegator is a type object that simply delegates all TypeMetadata calls to another TypeMetadata.
<u>TypeMetadataFactoryPluggableService</u>	The TypeMetadataFactoryPluggableService class is responsible for loading TypeMetadataFactories.
<u>TypeMetadataFactoryStore</u>	This Class stores TypeMetadataFactories.
<u>TypePreferences</u>	This class is used by TypeMetadata instances to pass information about preferred TypeMetadata objects.
<u>Types</u>	The Types class is a holder class for the RESOURCES global variable for resource properties of the com.xis.types package.
<u>URLTypeMetadata</u>	URLTypeMetadata is a type object that supports URL objects.

Exception Summary	
<u>NoSuchEnumerationException</u>	Class to implement an enumeration exception in Java.
<u>ParseFailedException</u>	A ParseFailedException is thrown (typically by TypeIO objects) when it is not possible to parse a given String as desired.
<u>TestFailedException</u>	The TestFailedException is thrown from the "test()" method of a DataTest subclass when the value fails the test.

Package com.xis.types Description

This package contains classes that provide several standard TypeMetadata classes for describing types and their constraints, and for rendering and editing values of those types.

FIG. 12D




Figure 13




Figure 14

Figure 15

HelloWorld.java

/* XIS Tutorial standalone sequence example 1 data class. */

```
public class HelloWorld {  
    private float value = 1.5f;  
  
    public String toString() {  
        return "Hello World!";  
    }  
  
    public int getID() {  
        return 5;  
    }  
  
    public float getValue() {  
        return value;  
    }  
  
    // uncomment this to make "value" editable  
    /*  
    public void setValue(float value) {  
        this.value = value;  
    }  
    */  
}
```

Figure 16A

TestHarness.java

```

/* XIS Tutorial standalone sequence example 1 XIS interfacing. */
import javax.swing.JFrame;
import java.awt.event.WindowAdapter;
import java.awt.event.WindowEvent;
import com.xis.propertysheet.PropertySheetInfoBean;
import com.xis.ui.UIBeanEvent;
import com.xis.ui.UIBeanAdapter;
import com.xis.leif.im.BaseInfoModel;
    } 1602

public class TestHarness {

    public static void main(String[] args) {

        // the plugin manager is only required for more complex applications
        // involving multiple components integrated at runtime
        BaseInfoModel.setStartingPluginManager(false);
    } 1604

        // a property sheet infobean to display HelloWorld's attributes
        PropertySheetInfoBean properties = new PropertySheetInfoBean();
        properties.addRowDataItem(new HelloWorld());
    } 1606

        // add listener for 'OK,' 'cancel,' or close, which generate 'close' events
        properties.addUIBeanListener(
            new UIBeanAdapter() {
                public void closed(UIBeanEvent event) {
                    System.exit(0);
                }
            }
        );
    } 1608

        // a top-level frame to hold our property sheet infobean
        JFrame frame = new JFrame("HelloWorld Properties");
        // add a listener for window closing
        frame.addWindowListener(
            new WindowAdapter() {
                public void windowClosing(WindowEvent e) {
                    System.exit(0);
                }
            }
        );
    } 1610A
};

```

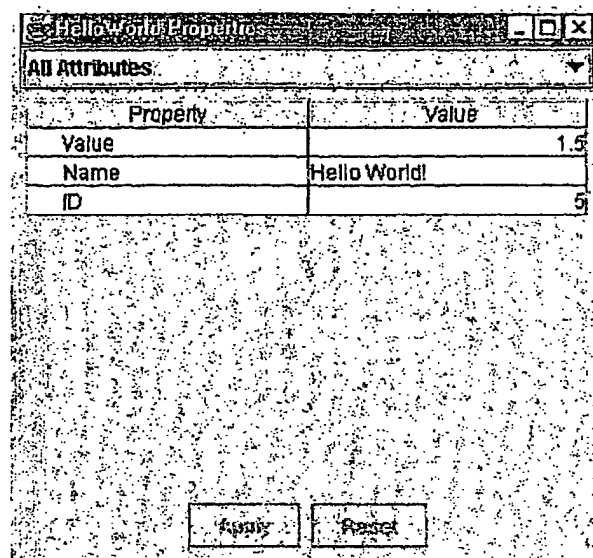
Figure 16B

Continuation of TestHarness.java

```
// stick the bean in the frame and display it
frame.getContentPane().add(properties);
frame.pack();
frame.setVisible(true);
}
}
```

} 1610B

Figure 17



The image shows a screenshot of a software window titled "Hello World Properties". Inside the window, there is a section labeled "All Attributes" with a dropdown arrow. Below this is a table with two columns: "Property" and "Value". The table contains three rows of data. At the bottom of the window, there are two buttons labeled "Apply" and "Reset".

Property	Value
Value	1.5
Name	Hello World!
ID	5

Figure 18A

HelloWorld.java

/ XIS Tutorial standalone sequence example 2 data class. */*

*/**

import java.awt.Color;

import java.beans.PropertyChangeSupport;

import java.beans.PropertyChangeListener;

**/*

public class HelloWorld {

*/**

private int value = 1;

private Color myColor = Color.green;

// this member class helps distribute property change events within XIS
private PropertyChangeSupport propertyChangeSupport =
new PropertyChangeSupport(this);

// two aux methods to let other XIS objects pay attention to this one
public void addPropertyChangeListener(PropertyChangeListener l) {
propertyChangeSupport.addPropertyChangeListener(l);
}

public void removePropertyChangeListener(PropertyChangeListener l) {
propertyChangeSupport.removePropertyChangeListener(l);
}

public String toString() {
return "A HelloWorld Object";
}

public String getGreeting() {
return "Hello World!";
}

**/*

public int getID() {
return 5;
}

1806

1802

Figure 18B

Continuation of of HelloWorld.java

```
public /*{*/ int /*}*/ getValue() {  
    return value;  
}
```

```
/*{*/  
1803 { public void setValue(int value) {  
    // only update and fire property change if this is really a change  
    if (this.value != value) {  
        int oldValue = this.value;  
        this.value = value;  
        // fire property change event to notify other XIS objects  
        propertyChangeSupport.firePropertyChange("value", oldValue, value);  
    }  
1804 }
```

```
public Color getMyColor() {  
    return myColor;  
}
```

```
1810 { public void setMyColor(Color myColor) {  
    // only update and fire property change if this is really a change  
    if (this.myColor != myColor) {  
        Color oldMyColor = this.myColor;  
        this.myColor = myColor;  
        // fire property change event to notify other XIS objects  
        propertyChangeSupport.firePropertyChange("myColor",  
1804         oldMyColor, myColor);  
    }  
1811 }  
/*}*/  
}
```

Figure 19A

TestHarness.java

```

/* XIS Tutorial standalone sequence step 2 XIS interfacing. */

import javax.swing.JFrame;
import java.awt.event.WindowAdapter;
import java.awt.event.WindowEvent;
import com.xis.propertysheet.PropertySheetInfoBean;
import com.xis.ui.UIBeanEvent;
import com.xis.ui.UIBeanAdapter;
import com.xis.leif.im.BaseInfoModel;
/*{*/
import jclass.chart.JCChart;
import com.xis.plot.PlotInfoBean;
import com.xis.plot.chartviews.LeifChartView;
/*}*/

public class TestHarness {

    public static void main(String[] args) {

        // the plugin manager is only required for more complex applications
        // involving multiple components integrated at runtime
        BaseInfoModel.setStartingPlugInManager(false);

        HelloWorld hello = new HelloWorld();

        // a property sheet infobean to display HelloWorld's attributes
        PropertySheetInfoBean properties = new PropertySheetInfoBean();
        properties.addRowDataItem(hello);

        // add a listener for 'OK' or 'cancel', which generate 'close' events
        properties.addUIBeanListener(
            new UIBeanAdapter() {
                public void closed(UIBeanEvent event) {
                    System.exit(0);
                }
            }
        );
    }
};

```

1902

Figure 19B

Continuation of TestHarness.java

```

/**
 // a top-level frame to hold our property sheet infobean
 JFrame propertySheetFrame = new JFrame("HelloWorld Properties");
 // add a listener for window closing
 propertySheetFrame.addWindowListener(
     new WindowAdapter() {
         public void windowClosing(WindowEvent e) {
             System.exit(0);
         }
     }
 );

 // stick the property sheet bean in the frame and display it
 propertySheetFrame.getContentPane().add(properties);
 propertySheetFrame.pack();
 propertySheetFrame.setVisible(true);

 // now we create a plot infobean to plot HelloWorld's numeric attribute
 PlotInfoBean plot = new PlotInfoBean();
 plot.addRowDataItems(new Object[] { hello });
 plot.setChartType(JCChart.BAR);
 // the alternatives are SCATTER_PLOT, PLOT, AREA, PIE, CANDLE,
 // and STACKING_BAR, though not all will make sense in this example

 // a top-level frame as before to hold our property sheet infobean
 JFrame plotFrame = new JFrame("HelloWorld Plot");

 // stick the plot bean in and put it up
 plotFrame.getContentPane().add(plot);
 plotFrame.pack();
 plotFrame.setVisible(true);
 */
}
}

```

1906 {

1904 }

Figure 20

The image shows a screenshot of a Windows-style dialog box titled "HelloWorld Properties". Inside the dialog, there is a tab labeled "All Attributes". Below the tab is a table with two columns: "Property" and "Value". The table contains five rows of data. At the bottom of the dialog, there are two buttons: "Apply" and "Reset".

Property	Value
ID	5
Name	A HelloWorld Object
Value	1
Greeting	Hello World!
My Color	<input type="checkbox"/> <input type="checkbox"/>

Figure 21

HelloWorld.java

```
/* XIS Tutorial standalone sequence step 3 data class. */
import java.awt.Color;
// (property change support moved to HelloWorldTranslator)

public class HelloWorld {
    private int value = 1;
    private Color myColor = Color.green;

    public String toString() {
        return "A HelloWorld Object";
    }

    public String getGreeting() {
        return "Hello World!";
    }

    public int getID() {
        return 5;
    }

    public int getValue() {
        return value;
    }

    /*
    public void setValue(int value) {
        // all the worrying about change events is moved to the translator,
        // so we just need to do the bare change operation (unless nonXIS
        // components need to listen to PropertyChanges)
        this.value = value;
    }
    */

    public Color getMyColor() {
        return myColor;
    }

    /*
    public void setMyColor(Color myColor) {
        // just need to set the value (see setValue())
        this.myColor = myColor;
    }
    */
}
```

Figure 22A

HelloWorldTranslator.java

/* XIS Tutorial standalone sequence step 3 data translator class. */

```

/*
import com.xis.leif.im.AttributeGetRequest;
import com.xis.leif.im.AttributeSetRequest;
import com.xis.leif.im.Domain;
import com.xis.leif.im.Translator;
import com.xis.leif.im.FieldMetaData;
import com.xis.domains.display.DisplayDomain;
import com.xis.domains.movement.MovementDomain;
import java.awt.Color;

```

```

public class HelloWorldTranslator extends Translator {

```

2202 {
 // the domains from which canned attribute metadata will be taken
 // NOTE, if an attribute appears in the methods below but its domain
 // is NOT listed here, THE ATTRIBUTE WILL BE IGNORED BY XIS
 private static final Domain[] baseDomains = new Domain[] {
 DisplayDomain.getDomain(), MovementDomain.getDomain()
 };

2203 {
 // store info about the fields, such as whether they are preferred or not
 private FieldMetaData[] fieldMetaDataArray;

2204 {
 // Return the Domains that describe the Attributes.
 public Domain[] getBaseDomains() {
 return baseDomains;
 }

// this method returns info on each field defined in the methods below
 public FieldMetaData[] getFieldMetaDataArray() {

2206 {
 if (fieldMetaDataArray == null) {
 // initialize default metadata
 FieldMetaData dispname = new
 FieldMetaData(DisplayDomain.displayName);
 FieldMetaData pencolor = new
 FieldMetaData(DisplayDomain.penColor);
 FieldMetaData speed = new
 FieldMetaData(MovementDomain.speed);
 FieldMetaData course = new
 FieldMetaData(MovementDomain.course);

Figure 22B

Continuation of HelloWorldTranslator.java

```

// attributes are visible ('preferred') by default; this
// turns this off for the course attribute
course.setVisibility(false);
// the order we put the attributes in here determines the order
// they appear in tables or property sheets
fieldMetaDataArray = new FieldMetaData[] {
    dispname, speed, course, pencolor
};
}

return fieldMetaDataArray;
}

////////////////////////////////////
// the following methods expose attributes of the HelloWorld class;
// instead of calling the class methods directly, XIS will access
// everything through this translator class
////////////////////////////////////
public String getDisplayName(AttributeGetRequest attributeGetRequest) {
    return ((HelloWorld)
        attributeGetRequest.getRawDataItem()).toString();
}

public Color getPenColor(AttributeGetRequest attributeGetRequest) {
    return ((HelloWorld)
        attributeGetRequest.getRawDataItem()).getMyColor();
}

public void setPenColor(AttributeSetRequest attributeSetRequest,
    Color penColor) {
    HelloWorld helloWorld = (HelloWorld)
        attributeSetRequest.getRawDataItem();
    Color oldPenColor = helloWorld.getMyColor();
    if (!penColor.equals(oldPenColor)) {
        helloWorld.setMyColor(penColor);
        // fire property change event to notify other XIS objects
        attributeSetRequest.getBaseDataItem().fireAttributeChanged(
            DisplayDomain.penColor, oldPenColor, penColor, true);
    }
}

public double getSpeed(AttributeGetRequest attributeGetRequest) {
    return (double) ((HelloWorld)
        attributeGetRequest.getRawDataItem()).getValue();
}

```

2206 B

2210

Figure 22C

Continuation of HelloWorldTranslator.java

```

2210 {
    public void setSpeed(AttributeSetRequest attributeSetRequest,
        double speed) {
        HelloWorld helloWorld = (HelloWorld)
            attributeSetRequest.getRawDataItem();
        Double oldSpeed = new Double((double)helloWorld.getValue());
        if (oldSpeed.doubleValue() != speed) {
            helloWorld.setValue((int)speed);
            // fire property change event to notify other XIS objects
            attributeSetRequest.getBaseDataItem().fireAttributeChanged(
                MovementDomain.speed, oldSpeed, new Double(speed), true);
        }
    }

    // this is a dummy attribute to demonstrate field metadata
    public double getCourse(AttributeGetRequest attributeGetRequest) {
        return (double) 0;
    }

    /**/
    // uncomment this to allow reflection to expose additional attributes
    // (see documentation under "Fooling Around")
    // public HelloWorldTranslator() {
    //     introspectExcept(new String[] {"value", "myColor"});
    // }
    /**/

}
/**/

```

2212 }

2208 }

Figure 23A

TestHarness.java

```

/* XIS Tutorial standalone sequence step 3 XIS interfacing. */

import javax.swing.JFrame;
import java.awt.event.WindowAdapter;
import java.awt.event.WindowEvent;
import com.xis.propertysheet.PropertySheetInfoBean;
import com.xis.ui.UIBeanEvent;
import com.xis.ui.UIBeanAdapter;
import com.xis.leif.im.BaseInfoModel;
import jclass.chart.JCChart;
import com.xis.plot.PlotInfoBean;
import com.xis.plot.chartviews.LeifChartView;
/*{*/
import com.xis.leif.im.TranslatorRegistry;
import com.xis.leif.im.LeifDataItem;
import com.xis.leif.im.InfoModel;
import com.xis.leif.im.BaseInfoModel;
import com.xis.domains.movement.MovementDomain;
import com.xis.domains.movement.MovementDomainWrapper;
import com.xis.leif.im.LeifDataItemDelegator;
import java.lang.reflect.InvocationTargetException;
import com.xis.leif.im.UndefinedLeifAttributeException;
/*}*/

public class TestHarness {

/*{*/
    protected static LeifDataItem leifHello;

    static {
        // Register the translator for HelloWorld. In fact this is really
        // only necessary when we have not followed the standard naming
        // convention (see docs), but it can't hurt.
        TranslatorRegistry.getTranslatorRegistry().registerObjectSchema(
            HelloWorld.class, HelloWorldTranslator.class);
    }
/*}*/

    public static void main(String[] args) {

        // the plugin manager is only required for more complex applications
        // involving multiple components integrated at runtime
        BaseInfoModel.setStartingPluginManager(false);

        HelloWorld hello = new HelloWorld();

```

} 2302

Figure 23B

Continuation of TestHarness.java

```
// a property sheet info bean to display HelloWorld's attributes
PropertySheetInfoBean properties = new PropertySheetInfoBean();
properties.addRawDataItem(hello);

// add a listener for 'OK' or 'cancel', which generate 'close' events
properties.addUIBeanListener(
    new UIBeanAdapter() {
        public void closed(UIBeanEvent event) {
            System.exit(0);
        }
    }
);

// a top-level frame to hold our property sheet info bean
JFrame propertySheetFrame = new JFrame("HelloWorld Properties");
// add a listener for window closing
propertySheetFrame.addWindowListener(
    new WindowAdapter() {
        public void windowClosing(WindowEvent e) {
            System.exit(0);
        }
    }
);

// stick the property Sheet bean in the frame and display it
propertySheetFrame.getContentPane().add(properties);
propertySheetFrame.pack();
propertySheetFrame.setVisible(true);
// now we create a plot info bean to plot HelloWorld's numeric attribute
PlotInfoBean plot = new PlotInfoBean();
plot.addRawDataItems(new Object[] { hello });
plot.setChartType(JCChart.BAR);
// the alternatives are SCATTER_PLOT, PLOT, AREA, PIE, CANDLE,
// and STACKING_BAR, though not all will make sense in this example

// We can set the attribute for initial display on the plot;
// if we do, this must consist of the attribute name preceded
// by the fully-qualified classname which ORIGINALLY DEFINES
// the attributeDescriptor — i.e., using "HelloWorld.speed"
// here will NOT work! If the descriptor is not defined in a
// domain or translator class, then it will have been defined
// dynamically through introspection when the first instance
// of the data item is dropped into an XIS InfoBean.
plot.setYAxisAttribute(
    "com.xis.domains.movement.MovementDomain.speed");
```

Figure 23C

Continuation of TestHarness.java

```

plot.setDynamicAdjustment(true); // so axes track value magnitude
plot.setBarChartAdjusting(true); // needed in some cases for bar chart

// a top-level frame as before to hold our plot info bean
JFrame plotFrame = new JFrame("HelloWorld Plot");

// stick the plot bean in and put it up
plotFrame.getContentPane().add(plot);
plotFrame.pack();
plotFrame.setVisible(true);

/*
// create a LeifDataItem version of hello and start a thread that
// will increase it
leifHello =
    BaseInfoModel.getBaseInfoModel().getLeifDataItem(hello);
new Accelerate();
*/
} // main

}

/*
// thread to update the speed attribute on the leifHello instance we created
class Accelerate extends Thread {

    public Accelerate() {
        super("Accelerator Thread");
        start();
    }

    public void run() {

        // wrap the LeifDataItem leifHello in a convenience wrapper that
        // gives access to attributes within that domain, if they exist
        MovementDomainWrapper helloMovementWrapper =
            MovementDomain.takeWrapper(TestHarness.leifHello);

```

Figure 23D

Continuation of TestHarness.java

```
while (true) {  
    // sleep for 0.5 seconds, then...  
    try {  
        sleep(500);  
    } catch (InterruptedException e) {  
        System.exit(1);  
    }  
  
    // ..update the speed attribute  
    try {  
        helloMovementWrapper.setSpeed(  
            helloMovementWrapper.getSpeed()+1);  
    } catch (UndefinedLeifAttributeException ulae) {  
        // exception if this data item doesn't have this attribute  
        System.exit(1); // usually we would do something better  
    } catch (InvocationTargetException ite) {  
        // sweep up any exception tossed by the underlying raw item  
        System.exit(1); // usually we would do something better  
    }  
} // while  
} // run()  
};  
/* */
```

Figure 24A

Property	Value
Name	A HelloWorld Object
Speed (KTS)	5
Course (DEG)	0
Pen Color	

Apply Reset

Figure 24B




Figure 25A

HelloWorld.java

/* XIS Tutorial standalone sequence step 4 data class. */

import java.awt.Color;

/*{*/

import com.xis.leif.im.FieldMetaData;

import com.xis.domains.display.DisplayDomain;

import com.xis.domains.movement.MovementDomain;

import com.xis.leif.im.Domain;

import com.xis.leif.im.AttributeGetRequest;

import com.xis.leif.im.AttributeSetRequest;

import com.xis.leif.im.AttributeDescriptor;

import java.beans.PropertyChangeSupport;

import java.beans.PropertyChangeListener;

/*}*/

public class HelloWorld {

private int value = 1;

private Color myColor = Color.green;

/*{*/

public static AttributeDescriptor getDisplayNameDescriptor() {
return DisplayDomain.displayName;

}

public static AttributeDescriptor getSpeedDescriptor() {
return MovementDomain.speed;

}

public static AttributeDescriptor getPenColorDescriptor() {
return DisplayDomain.penColor;

}

/*}*/

/*{*/ // this property change support code as in step 2 /*}*/

// this member class helps distribute property change events within XIS
private PropertyChangeSupport propertyChangeSupport =
new PropertyChangeSupport(this);

// two aux methods to let other XIS objects pay attention to this one

public void addPropertyChangeListener(PropertyChangeListener l) {
propertyChangeSupport.addPropertyChangeListener(l);

}

public void removePropertyChangeListener(PropertyChangeListener l) {
propertyChangeSupport.removePropertyChangeListener(l);

}

2502

2504

Figure 25B

Continuation of HelloWorld.java

```
public String toString() {
    return "A HelloWorld Object";
}

public String getGreeting() {
    return "Hello World!";
}

public int getID() {
    return 5;
}

public int getValue() {
    return value;
}

/*{! // this function as in step 2 !}{*/
public void setValue(Int value) {
    // we only want to update and fire property change if really changes
    if (this.value != value) {
        int oldValue = this.value;
        this.value = value;
        // fire property change event to notify other XIS objects
        propertyChangeSupport.firePropertyChange("value", oldValue, value);
    }
}

/*{*/
// "myColor"-related methods changed to expose "penColor" instead

public Color getPenColor() {
    return myColor;
}

public void setPenColor(Color penColor) {
    // we only want to update and fire property change if really changes
    if (penColor != this.myColor) {
        Color oldPenColor = this.myColor;
        this.myColor = penColor;
        // fire property change event to notify other XIS objects
        propertyChangeSupport.firePropertyChange("penColor",
            oldPenColor, penColor);
    }
}
```

Figure 25C

Continuation of HelloWorld.java

```

// expose toString() return under a new name
public String getDisplayName() {
    return toString();
}
// expose "value" under a new name
public double getSpeed() {
    return (double) getValue();
}

public void setSpeed(double speed) {
    // we only want to update and fire property change if really changes
    Double oldSpeed = new Double((double)this.getValue());
    if (speed != oldSpeed.doubleValue()) {
        setValue((int)speed);
        // fire property change event to notify other XIS objects
        propertyChangeSupport.firePropertyChange("speed", oldSpeed,
            new Double(speed));
    }
}

/**
// store info about the fields, such as whether they are preferred or not
private static FieldMetaData[] fieldMetaDataArray;

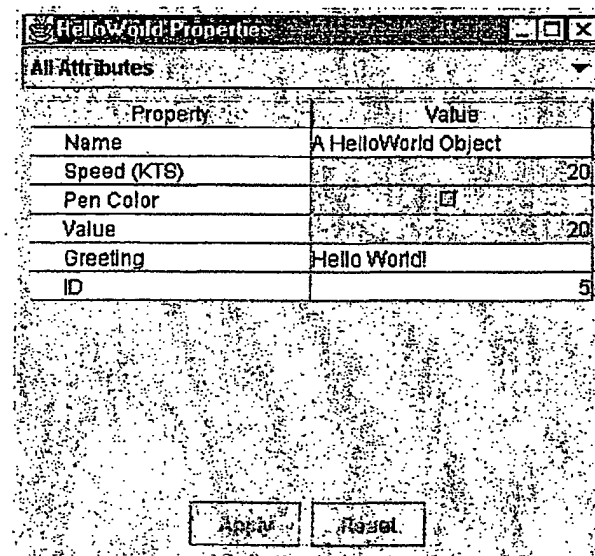
// this method returns info on each field defined in the methods below
public static FieldMetaData[] getFieldMetaDataArray() {

    if (fieldMetaDataArray == null) {
        // initialize default metadata
        FieldMetaData dispname = new
            FieldMetaData(DisplayDomain.displayName);
        FieldMetaData pencolor = new
            FieldMetaData(DisplayDomain.penColor);
        FieldMetaData speed = new
            FieldMetaData(MovementDomain.speed);
        // could customize the field metadata here
        fieldMetaDataArray = new FieldMetaData[] {
            dispname, speed, pencolor
        };
    }
    return fieldMetaDataArray;
}
*/
}

```

2506

Figure 26



The screenshot shows a window titled "HelloWorld Properties" with a standard Windows-style title bar (minimize, maximize, close buttons). Below the title bar is a dropdown menu currently set to "All Attributes". The main area of the window contains a table with two columns: "Property" and "Value". The table lists several attributes of a HelloWorld object. At the bottom of the window, there are two buttons: "Apply" and "Reset".

Property	Value
Name	A HelloWorld Object
Speed (KTS)	20
Pen Color	<input checked="" type="checkbox"/>
Value	20
Greeting	Hello World!
ID	5

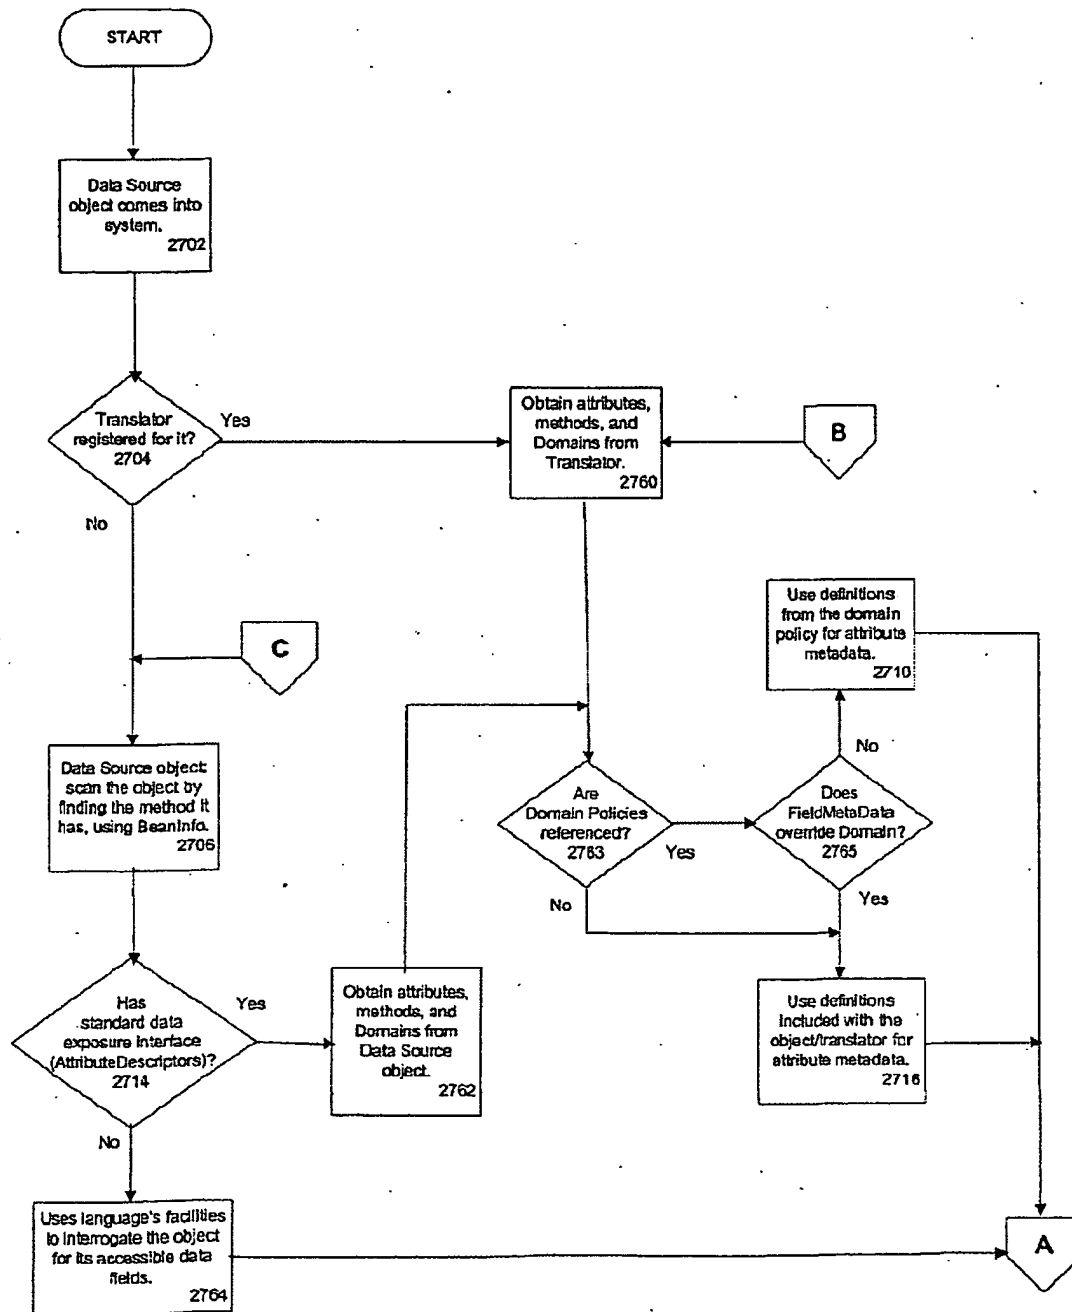


Figure 27A

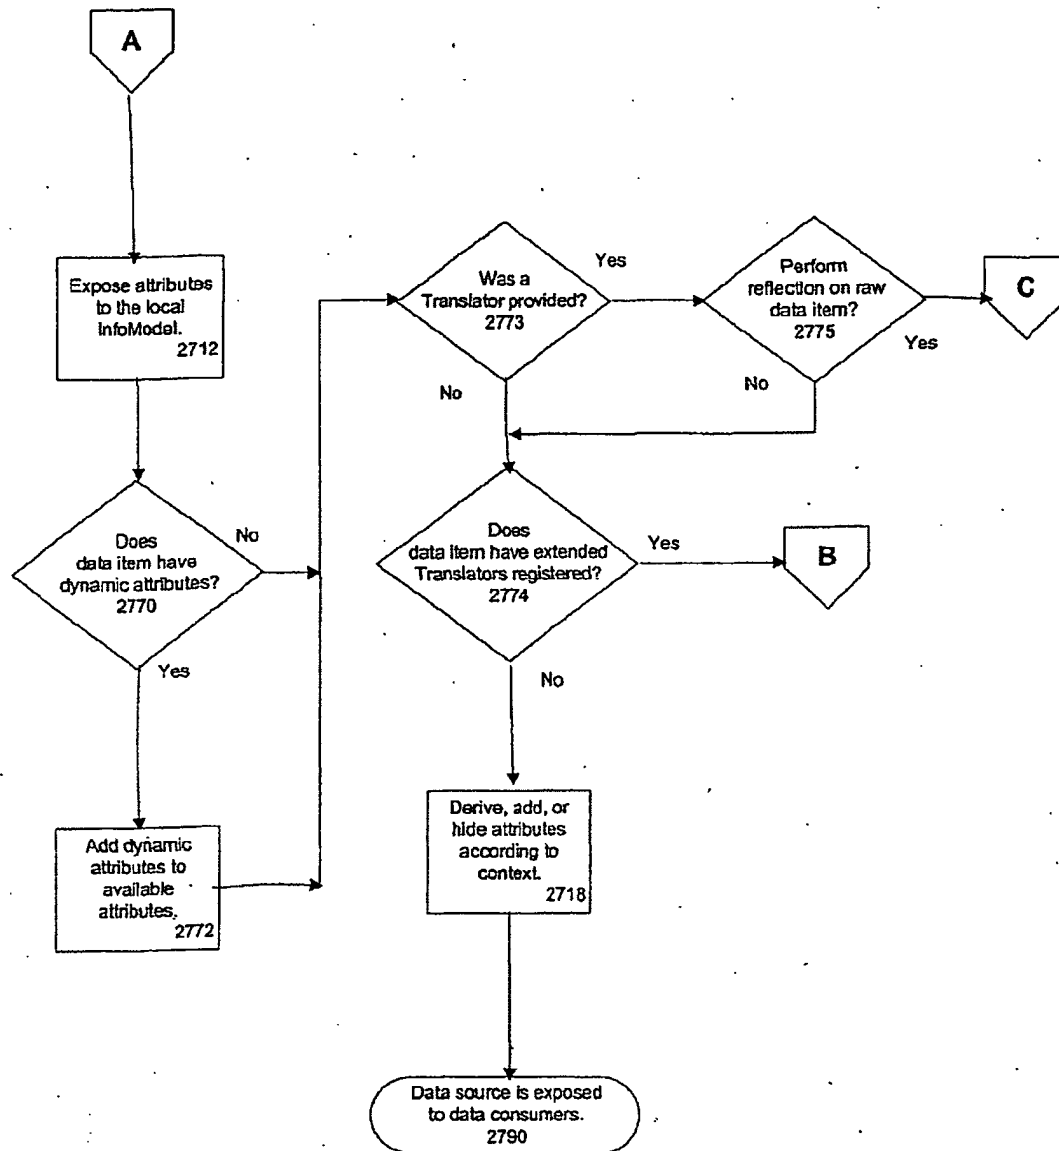


Figure 27B




Figure 28

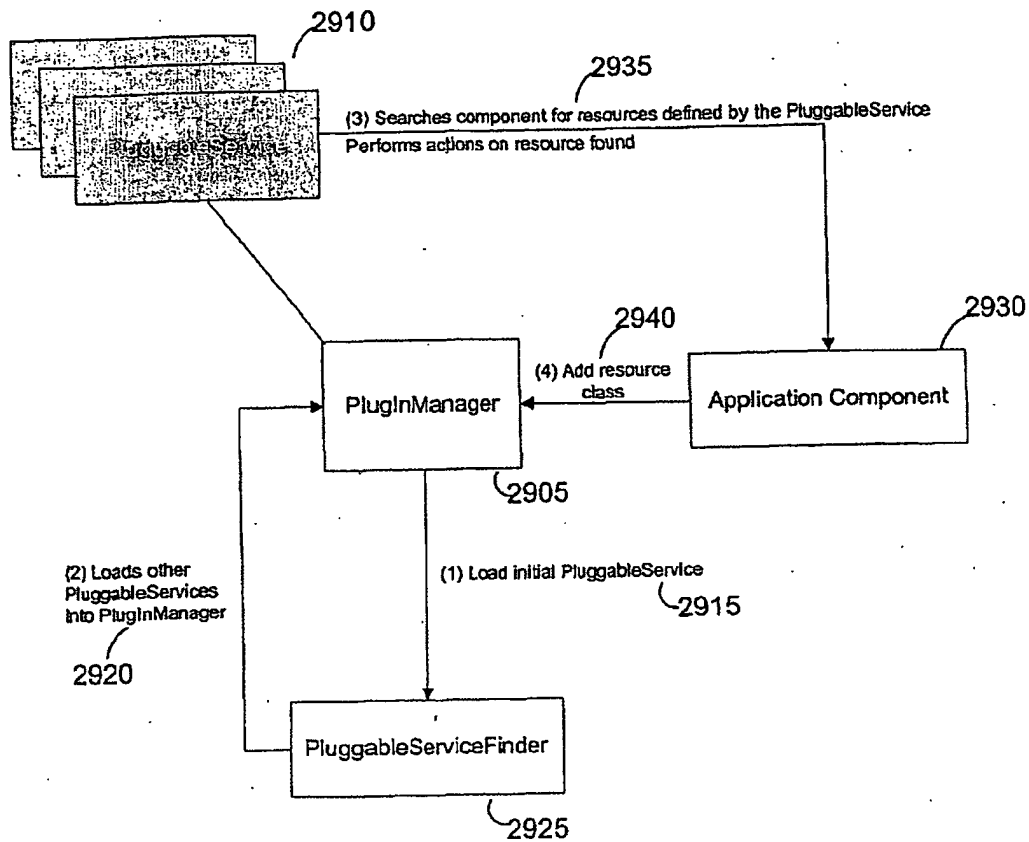


Figure 29

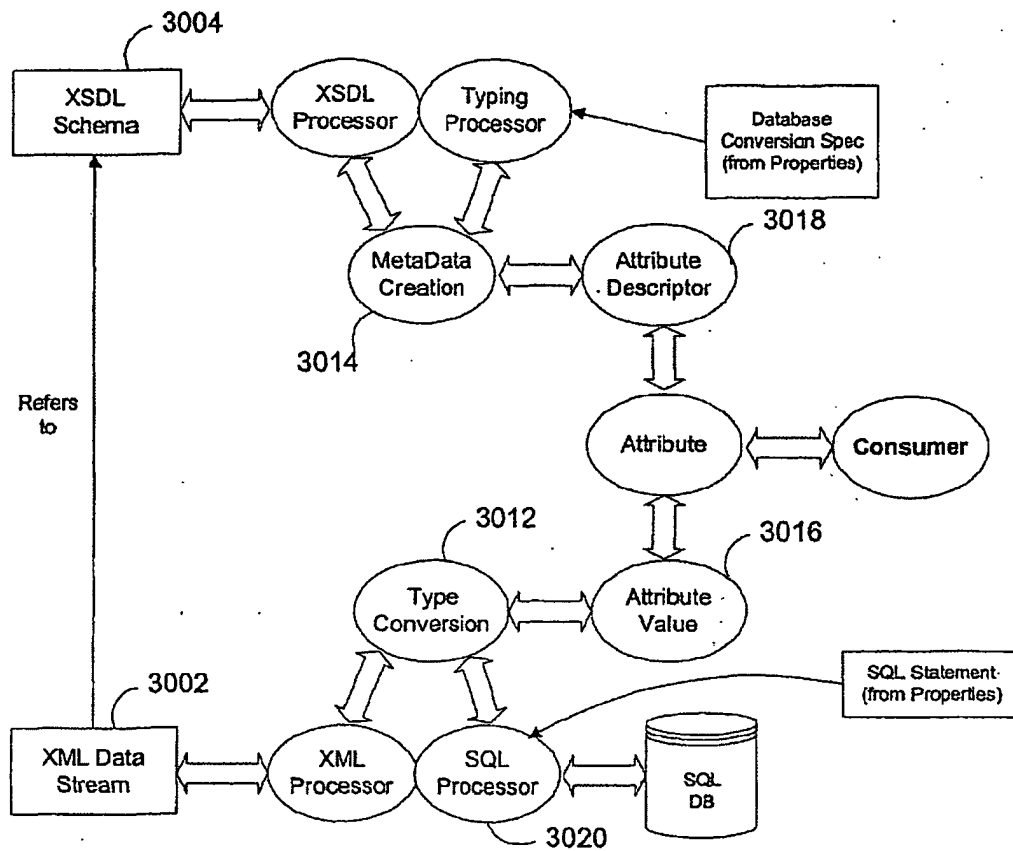


Figure 30

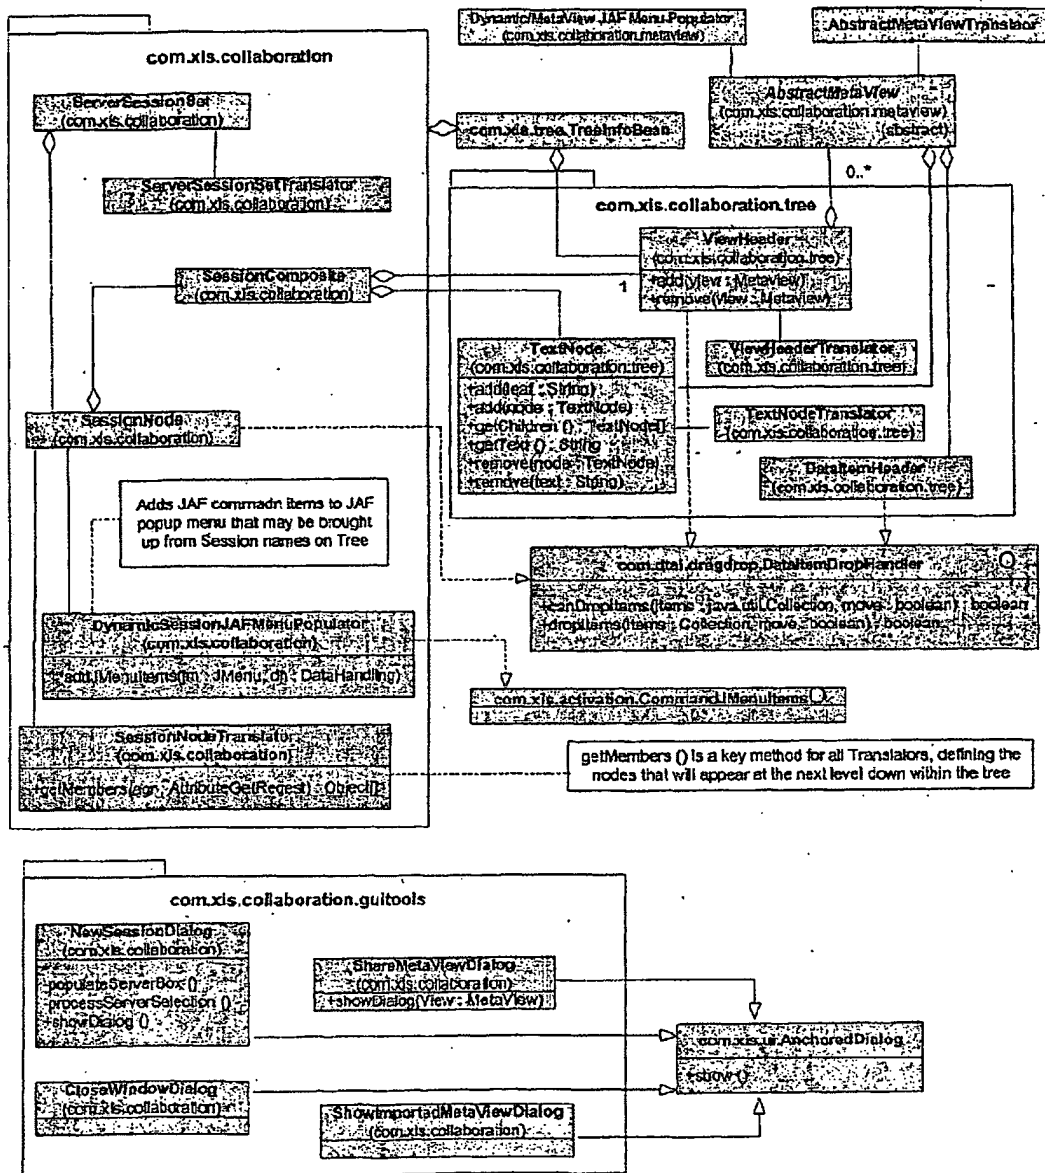


Figure 31

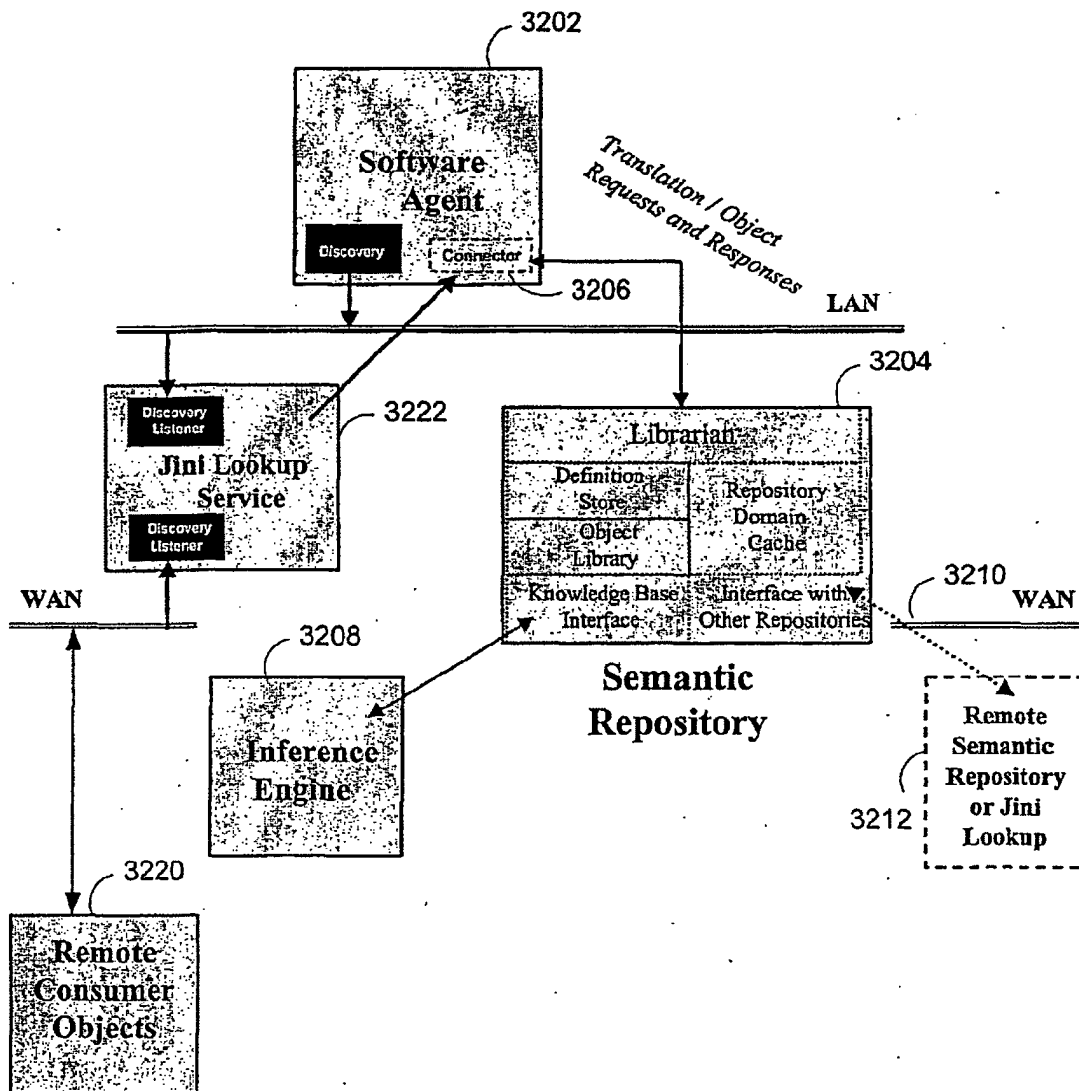


Figure 32

Figure 33A

Class ContentInfoBean

```

java.lang.Object
|-- java.awt.Component
    |-- java.awt.Container
        |-- javax.swing.JComponent
            |-- javax.swing.JPanel
                |-- com.xis.ui.AbstractUIBean
                    |-- com.xis.leif.infobeans.DataItemSinkUIBean
                        |-- com.xis.infobeans.content.ContentInfoBean

```

All Implemented Interfaces:

Accessible, BeanContextChildOwner, BeanContextChildOwnerDelegator, BeanContextProxy,
BeanContextServicesOwnerDelegator, ClipboardUser, DataItemSink, ImageObserver, MenuContainer,
Serializable, StateSavable, UIBean

```

public class ContentInfoBean
extends DataItemSinkUIBean
implements ClipboardUser

```

The ContentInfoBean class is a visual component that displays the contents of a raw data item. If no contents are available, it defaults to a split pane containing the JAF menu and the PropertySheet of the raw data item. The contents may have be multipart, and may be text, html, rich text, or an image. Multimedia support will soon be added.

Author:

Jaime Garcia, Polaxis, Inc.

See Also:

Serialized Form

Inner classes inherited from class javax.swing.JPanel

JPanel.AccessibleJPanel

Inner classes inherited from class javax.swing.JComponent

JComponent.AccessibleJComponent

Inner classes inherited from class java.awt.Container

Container.AccessibleAWTContainer

Figure 33B

Inner classes inherited from class java.awt.Component	
<u>Component.AccessibleAWTComponent</u>	
Field Summary	
static <u>URLConnectionProvider</u>	RESOURCES localized resources for this view object.
Fields inherited from class com.xis.ui.AbstractUIBean	
<u>uiBeanListener</u>	
Fields inherited from class javax.swing.JComponent	
<u>accessibleContext</u> , <u>listenerList</u> , <u>TOOL_TIP_TEXT_KEY</u> , <u>ui</u> , <u>UNDEFINED_CONDITION</u> , <u>WHEN_ANCESTOR_OF_FOCUSED_COMPONENT</u> , <u>WHEN_FOCUSED</u> , <u>WHEN_IN_FOCUSED_WINDOW</u>	
Fields inherited from class java.awt.Component	
<u>BOTTOM_ALIGNMENT</u> , <u>CENTER_ALIGNMENT</u> , <u>LEFT_ALIGNMENT</u> , <u>RIGHT_ALIGNMENT</u> , <u>TOP_ALIGNMENT</u>	
Fields inherited from interface java.awt.image.ImageObserver	
<u>ABORT</u> , <u>ALLBITS</u> , <u>ERROR</u> , <u>FRAMEBITS</u> , <u>HEIGHT</u> , <u>PROPERTIES</u> , <u>SOMEBITS</u> , <u>WIDTH</u>	
Constructor Summary	
<u>ContentInfoBean()</u> Default constructor that creates an empty ContentInfoBean.	
Method Summary	
void	<u>addRawDataItem</u> (Object rawDataItem) Load the raw data item into the ContentInfoBean.
void	<u>addRawDataItems</u> (Object[] rawDataItems) Add the raw data items in the array.
boolean	<u>canClear</u> () Return true if the ContentInfoBean has an object and it is selected
boolean	<u>canClear</u> (Object[] items) Return true if the specified items can be cleared.
boolean	<u>canCopy</u> () Return true if the ContentInfoBean has an object and it is selected
boolean	<u>canCopy</u> (Object[] items) Return true if the specified items can be copied.
boolean	<u>canCut</u> () Return true if the ContentInfoBean has an object and it is selected
boolean	<u>canCut</u> (Object[] items) Return true if the specified items can be cut.
boolean	<u>canPaste</u> () Return true if the ContentInfoBean can paste new objects, false if not.
boolean	<u>canSelectAll</u> () Return true if the ContentInfoBean can select all objects.

Figure 33c

boolean	<u>canSelectNone()</u> Return true if the ContentInfoBean can un-select all objects.
void	<u>clear()</u> Notify the ContentInfoBean to remove the current raw data item only if it is selected.
void	<u>clear(Object[] items)</u> Clears the given items.
void	<u>clearAll()</u> Removes the currently loaded object.
boolean	<u>contains(Object[] items)</u> Return true if this ContentInfoBean contains <i>all</i> the objects of the given array.
boolean	<u>containsComponent(Component component)</u> Check if the given component is contained by this InfoBean
void	<u>copy(Clipboard clipboard)</u> Called to invoke this ContentInfoBean's copy action, which is to copy all selected data to the Clipboard.
void	<u>copy(Clipboard clipboard, Object[] items)</u> Copies the given items into the Clipboard.
protected void	<u>createContent()</u> Method called when there is no content to display for a raw data item.
void	<u>cut(Clipboard clipboard)</u> Cut selected items from the ContentInfoBean and post them into Clipboard.
void	<u>cut(Clipboard clipboard, Object[] items)</u> Cut the given items from the ContentInfoBean and post them into the given Clipboard only if they occur in the ContentInfoBean.
protected Container	<u>getContainerForContent(int index)</u> Get a container with the contents of the content object at the given index, or null if the content type is not supported.
Object	<u>getContent()</u> Fetch the currently loaded raw data item
JAFAndPropertyComponent	<u>getJAFAndPropertyComponent()</u> Get the JAFAndProperty component used by the ContentInfoBean to display the contents for raw data items that have nothing else to display.
Menu	<u>getLeifDataItemMenu(LeifDataItem dataItem, boolean showCutPasteItems)</u> Return the data item menu for a LeifDataItem (usually the selected LeifDataItem).
TypedResourceBundle	<u>getResources()</u> Return the ResourceBundle for this ContentInfoBean.
Object[]	<u>getSelectedObjects()</u> Get an array of selected objects.
void	<u>infoModelChanged()</u> Messaged to indicate an infoModel change for this InfoBean or one or more of its LeifDataItems.
boolean	<u>isCreatingContent()</u> Check whether default content creation is set.
boolean	<u>isDragEnabled()</u> Return true if the default Drag support is enabled.
boolean	<u>isDropEnabled()</u> Return true if the default Drop support is enabled.

Figure 3

boolean	<code>isSelected()</code> Check the selected state of the content object, if there is currently one loaded.
boolean	<code>isXISNotifying()</code> Check whether the ContentInfoBean is updating based on XIS events and is notifying XIS of raw data item attribute changes.
void	<code>paste(Clipboard clipboard)</code> Paste the data Objects from the given clipboard.
void	<code>removeAllRawDataItems()</code> Remove all of the raw data items that are currently loaded.
void	<code>removeRawDataItem(Object rawDataItem)</code> Remove the given raw data item if it is the currently loaded raw data item.
void	<code>removeRawDataItems(Object[] rawDataItems)</code> Remove the raw data items in the array.
void	<code>selectAll()</code> Set the selection state of the object to true.
void	<code>selectNone()</code> Set the selection state of the object to false.
void	<code>setCreateContent(boolean create)</code> Set whether content should be created for objects that do not have any displayable content, via a JAFAndPropertyComponent.
void	<code>setDragEnabled(boolean enabledrag)</code> Set the status of the default Drag support.
void	<code>setDragOwnerProxy(DragOwner dragProxy)</code> Set a DragOwner "proxy" for this InfoBean.
void	<code>setDropEnabled(boolean enabledrop)</code> Set the status of the default Drop support.
void	<code>setDropOwnerProxy(DropOwner dropProxy)</code> Set a DropOwner "proxy" for this InfoBean.
void	<code>setSelection(boolean selected)</code> Set the selection state of the content object.
void	<code>setXISNotifying(boolean notify)</code> Set whether the ContentInfoBean should update based on XIS events and should notify XIS of raw data item attribute changes.

Methods inherited from class `com.xis.jei.infobean.DataItemSinkUIBean`

`addJAFPopulator, addRawDataItemsAsGroup, addService, close, createBeanContextServicesOwnerDelegator, dispose, getBeanContextProxy, getBeanContextServices, getEzContext, getInfoModel, getJAFPopulators, getLeafDataItemMenu, getLeafDataItemMenu, getLeafDataItemMenu, getMenuBar, getOwnedBeanContextChild, getService, getService, getService, getStatusBars, getToolBars, getUIs, initializeBeanContextResources, initializeBeanContextServices, invalidateInfoModel, isDropPasteEnabled, isHandlingClipboardOperations, releaseBeanContextResources, releaseBeanContextServices, revokeAnchoredDialogProvider, revokeFrameProvider, revokeService, setDropPasteEnabled, setHandlingClipboardOperations, validatePendingSetBeanContext`

Methods inherited from class `com.xis.ui.AbstractUIBean`

`addUIBeanListeners, finalize, getShortTitle, getTitle, getUIComponents, isActive, isClosed, isClosedOK, processUIBeanEvent, removeUIBeanListener, restoreState, saveState, setActive, setShortTitle, setTitle`

Figure 33F

Methods inherited from class java.lang.Object
clone, equals, getClass, hashCode, notify, notifyAll, wait, wait, wait

Methods inherited from interface com.xis.ui.ClipboardUser
setHandlingClipboardOperations, setHandlingClipboardOperations

Methods inherited from interface com.xis.ui.UIReg
addPropertyChangeListener, removePropertyChangeListener

Field Detail

RESOURCES

```
public static final TypedResourceBundle RESOURCES
```

localized resources for this view object.

Constructor Detail

ContentInfoBean

```
public ContentInfoBean()
```

Default constructor that creates an empty ContentInfoBean.

Method Detail

getContents

```
public Object getContents()
```

Fetch the currently loaded raw data item

Returns:

the currently loaded object, or null if no object is loaded

setXISNotifying

```
public void setXISNotifying(boolean notify)
```

Set whether the ContentInfoBean should update based on XIS events and should notify XIS of raw data item attribute changes

Parameters:

notify - if true then notify XIS, else do not

56/90

*Figure 29***isXISNotifying**

```
public boolean isXISNotifying()
```

Check whether the ContentInfoBean is updating based on XIS events and is notifying XIS of raw data item attribute changes

Returns:
true if it is notifying XIS, else false

removeAllRawDataItems

```
public void removeAllRawDataItems()
```

Remove all of the raw data items that are currently loaded

getResources

```
public TypedResourceBundle getResources()
```

Return the ResourceBundle for this ContentInfoBean.

Overrides:

getResources in class DataItemSinkUIBean

Returns:
the statically sourced ResourceBundle.

infoModelChanged

```
public void infoModelChanged()
```

Messaged to indicate an InfoModel change for this InfoBean or one or more of its LeftDataItems. This should reload all currently loaded data items to pick up InfoModel changes.

Overrides:

infoModelChanged in class DataItemSinkUIBean

addRawDataItems

```
public void addRawDataItems(Object[] rawDataItems)
```

Add the raw data items in the array. Will only add if the array has only one object and the object is not the currently loaded rawDataItem.

Parameters:

rawDataItems - the array of objects to be added

addRawDataItem

```
public void addRawDataItem(Object rawDataItem)
```

Figure 33H

Load the raw data item into the ContentInfoBean. If the raw data item supports the ContentDomain and has content of type that is currently supported then it will be displayed. Otherwise, if content creation is set to true then default content will be created, using a JAFAndPropertyComponent

Overrides:

addRawDataItem in class DataItemSinkUIBean

Parameters:

rawDataItem - the raw object to display

createContent

protected void createContent()

Method called when there is no content to display for a raw data item. The default behavior displays a JAFAndPropertyComponent of the object, but subclasses may wish to display something else instead.

isCreatingContent

public boolean isCreatingContent()

Check whether default content creation is set. If so then a JAFAndPropertyComponent will be created for items that have no displayable content.

Returns:

true if default content is created, else false

setCreateContent

public void setCreateContent(boolean create)

Set whether content should be created for objects that do not have any displayable content, via a JAFAndPropertyComponent.

Parameters:

create - if true, create content, else do not

removeRawDataItems

public void removeRawDataItems(Object[] rawDataItems)

Remove the raw data items in the array. Will only remove the object that is currently loaded if it is contained in the array

Parameters:

rawDataItems - the array of objects to be removed

removeRawDataItem

public void removeRawDataItem(Object rawDataItem)

Figure 33

Remove the given raw data item if it is the currently loaded raw data item

Overrides:

`removeRawDataItem` in class `DataItemSinkUIBean`

Parameters:

`rawDataItem` - the object to remove

`getJAFAndPropertyComponent`

`public JAFAndPropertyComponent getJAFAndPropertyComponent()`

Get the JAFAndProperty component used by the ContentInfoBean to display the contents for raw data items that have nothing else to display.

Returns:

the current JAFAndPropertyComponent being used

`getContainerForContent`

`protected Container getContainerForContent(int index)`

Get a container with the contents of the content object at the given index, or null if the content type is not supported. Subclass this if you need support for a type that is not already supported.

Parameters:

`index` - the index of in the object

Returns:

a Container with the contents at the given index

`getSelectedObjects`

`public Object[] getSelectedObjects()`

Get an array of selected objects. This will return an empty array if there are no selected objects. If the raw data item is selected it will return an array of size 1 with the raw data item inside.

Returns:

an array of selected objects

`isSelected`

`public boolean isSelected()`

Check the selected state of the content object, if there is currently one loaded. If there is none, return false.

Returns:

true if there is an object and it is selected

`selectAll`

`public void selectAll()`

Figure 30
PE

Set the selection state of the object to true.

Specified by:

selectAll in interface ClipboardUser

Overrides:

selectAll in class DataItemSinkUIBean

selectNone

public void selectNone()

Set the selection state of the object to false.

setSelection

public void setSelection(boolean selected)

Set the selection state of the content object

Parameters:

the --new selection state

canClear

public boolean canClear()

Return true if the ContentInfoBean has an object and it is selected

Specified by:

canClear in interface ClipboardUser

Overrides:

canClear in class DataItemSinkUIBean

Returns:

true if the raw data item is selected

See Also:

getSelectedObjects()

canClear

public boolean canClear(Object[] items)

Return true if the specified items can be cleared. If the item is not in the ContentInfoBean return false. If there is more than one item return false.

Specified by:

canClear in interface ClipboardUser

Overrides:

canClear in class DataItemSinkUIBean

Parameters:

items - the Array of items to be cut.

Returns:

true if all of the items are present, otherwise false.

See Also:

Figures 33-35

contains(Object[])

canCopy

public boolean canCopy()

Return true if the ContentInfoBean has an object and it is selected

Specified by:

canCopy in interface ClipboardUser

Overrides:

canCopy in class DataItemSinkUIBean

Returns:

true if the raw data item is selected

See Also:

canClear()

canCopy

public boolean canCopy(Object[] items)

Return true if the specified items can be copied. If the item is not in the ContentInfoBean return false. If there is more than one item return false.

Specified by:

canCopy in interface ClipboardUser

Overrides:

canCopy in class DataItemSinkUIBean

Parameters:

items - the Array of items to be cut.

Returns:

true if all of the items are present, otherwise false.

See Also:

canClear(Object[])

canCut

public boolean canCut()

Return true if the ContentInfoBean has an object and it is selected

Specified by:

canCut in interface ClipboardUser

Overrides:

canCut in class DataItemSinkUIBean

Returns:

true if the raw data item is selected

See Also:

canClear()

canCut

public boolean canCut(Object[] items)

Figure 33d

Return true if the specified items can be cut. If the item is not in the ContentInfoBean return false. If there is more than one item return false.

Specified by:

canCut in interface ClipboardUser

Overrides:

canCut in class DataItemSinkUIBean

Parameters:

items - the Array of items to be cut.

Returns:

true if all of the items are present, otherwise false.

See Also:

canClear(Object[])

canPaste

public boolean canPaste()

Return true if the ContentInfoBean can paste new objects, false if not.

Specified by:

canPaste in interface ClipboardUser

Overrides:

canPaste in class DataItemSinkUIBean

Returns:

default is always true for the ContentInfoBean.

canSelectAll

public boolean canSelectAll()

Return true if the ContentInfoBean can select all objects.

Specified by:

canSelectAll in interface ClipboardUser

Overrides:

canSelectAll in class DataItemSinkUIBean

Returns:

returns true if there is an object and it is not selected

canSelectNone

public boolean canSelectNone()

Return true if the ContentInfoBean can un-select all objects.

Returns:

returns true if the raw data item is selected

clear

public void clear()

62/90

Figure 3A4

Notify the ContentInfoBean to remove the current raw data item only if it is selected.

Specified by:

clear in interface ClipboardUser

Overrides:

clear in class DataItemsSinkUIBean

clear

public void clear(Object() items)

Cleares the given items. These items only get cleared if they actually occur in the ContentInfoBean.

Specified by:

clear in interface ClipboardUser

Overrides:

clear in class DataItemsSinkUIBean

Parameters:

items - the items to cleared.

clearAll

public void clearAll()

Removes the currently loaded object.

copy

public void copy(Clipboard clipboard)

Called to invoke this ContentInfoBean's copy action, which is to copy all selected data to the Clipboard.

Specified by:

copy in interface ClipboardUser

Overrides:

copy in class DataItemsSinkUIBean

Parameters:

clipboard - the Clipboard object that gets posted to. The actual items posted are contained in a LeifTransferable.

copy

public void copy(Clipboard clipboard,
 Object() items)

Copies the given items into the Clipboard.

Specified by:

copy in interface ClipboardUser

Overrides:

copy in class DataItemsSinkUIBean

Parameters:

clipboard - the Clipboard object that gets posted to. The actual items posted are contained in a

Figure 33 M
PCT

LeifTransferable.
items - the array of Object items to copied.

cut

public void cut(Clipboard clipboard)

Cut selected items from the ContentInfoBean and post them into Clipboard.

Specified by:

cut in interface ClipboardUser

Overrides:

cut in class DataItemSinkUIBean

Parameters:

clipboard - the Clipboard object that gets posted to. The actual items posted are contained in a LeifTransferable.

cut

public void cut(Clipboard clipboard,
Object[] items)

Cut the given items from the ContentInfoBean and post them into the given Clipboard only if they occur in the ContentInfoBean.

Specified by:

cut in interface ClipboardUser

Overrides:

cut in class DataItemSinkUIBean

Parameters:

clipboard - the Clipboard object that gets posted to. The actual items posted are contained in a LeifTransferable.

items - the array of Object items to cut and posted.

paste

public void paste(Clipboard clipboard)

Paste the data Objects from the given clipboard. This retrieves the clipboard contents and does a simple add.

Specified by:

paste in interface ClipboardUser

Overrides:

paste in class DataItemSinkUIBean

Parameters:

clipboard - the Clipboard that contains the objects.

See Also:

addRawDataItems(Object[])

contains

public final boolean contains(Object[] items)

Figure 33
PCT

Return true if this ContentInfoBean contains *all* the objects of the given array.

Parameters:

items - an array of objects to locate in the ContentInfoBean.

Returns:

true if there is only one object and it is contained, false otherwise

containsComponent

public boolean containsComponent(Component component)

Check if the given component is contained by this InfoBean

Parameters:

component - the Component to check for

Returns:

true if it is contained, else false

getLeifDataItemMenu

public JMenu getLeifDataItemMenu(LeifDataItem dataItem,
boolean showCutPasteItems)

Return the data item menu for a LeifDataItem (usually the selected LeifDataItem). This is used by the menu bar to add menu items from this JMenu to a Data menu if there is exactly one DataItemSelected selected.

Overrides:

getLeifDataItemMenu in class DataItemSinkUIBean

Parameters:

dataItem - the LeifDataItem to get the menu for

showCutPasteItems - true to allow cut and paste items to appear, false to omit them.

Returns:

the JMenu for the given LeifDataItem

isDragEnabled

public boolean isDragEnabled()

Return true if the default Drag support is enabled.

Overrides:

isDragEnabled in class DataItemSinkUIBean

Returns:

true if drag is enabled, false if not, default is initialized to true.

setDragEnabled

public void setDragEnabled(boolean enabledrag)

Set the status of the default Drag support.

Overrides:

setDragEnabled in class DataItemSinkUIBean

Parameters:

Figure 33

enabledrag - true if default drag support should be used, false if not.

setDragOwnerProxy

public void setDragOwnerProxy(DragOwner dragProxy)

Set a DragOwner "proxy" for this InfoBean.

Overrides:

setDragOwnerProxy in class DataItemSinkUIBean

Parameters:

dragProxy - a DragOwner implementation.

isDropEnabled

public boolean isDropEnabled()

Return true if the default Drop support is enabled.

Overrides:

isDropEnabled in class DataItemSinkUIBean

Returns:

true if drag is enabled, false if not, default is initialized to true.

setDropEnabled

public void setDropEnabled(boolean enabledrop)

Set the status of the default Drop support.

Overrides:

setDropEnabled in class DataItemSinkUIBean

Parameters:

enabledrop - true if default drag support should be used, false if not.

setDropOwnerProxy

public void setDropOwnerProxy(DropOwner dropProxy)

Set a DropOwner "proxy" for this InfoBean.

Overrides:

setDropOwnerProxy in class DataItemSinkUIBean

Parameters:

dropProxy - a DropOwner implementation.

Figure 3A

Package com.xis.leif.im

This package contains classes that provide the APIs for using information management in applications.

See:

Description

Interface Summary	
<u>Attribute</u>	The Attribute class represents an attribute for a particular data type.
<u>AttributeAlias</u>	The AttributeAlias indicates an alias from many Attributes to a single Attribute, together with a precision level; the higher the precision, the better the alias.
<u>AttributeFactory</u>	The AttributeFactory class allows an implementor to return an appropriate Attribute for the given LeifDataItem.
<u>AttributeLookup</u>	The AttributeLookup interface is used to lookup Attribute objects for a particular data item.
<u>DisplayLabel</u>	The DisplayLabel interface defines methods that are needed for use with DisplayLabelAttributes.
<u>Domain</u>	This interface describes the basic fields and methods possessed by all Domains.
<u>InfoModel</u>	The InfoModel interface is the interface that is used to convert raw data items into LeifDataItems.
<u>LeifDataItem</u>	The LeifDataItem interface represents a simple data item.
<u>LeifDataItemObserver</u>	This class is used for observing a LeifDataItem to know when it has finished processing an action.
<u>LiteDataItem</u>	The LiteDataItem interface represents a data item.
<u>PropertyProvider</u>	If a PropertyProvider implementation is added to services it can be used to replace the standard behavior when a PropertySheetView is opened from a JAF menu or as a default command.
<u>RawDataItemLookup</u>	The RawDataItemLookup interface is used to look up a raw data item from a unique id.

Class Summary	
<u>AbstractAttribute</u>	The AbstractAttribute class represents an Attribute.
<u>AttributeAliasPluggableService</u>	This registers AttributeAliases.
<u>AttributeDescriptor</u>	The AttributeDescriptor class is used to describe an attribute without providing functionality of how to use the attribute.
<u>AttributeDescriptorFactory</u>	The AttributeDescriptorFactory class is a singleton class used to create or get AttributeDescriptors.
<u>AttributeFactoryInfoModelSubset</u>	The AttributeFactoryInfoModelSubset class provides an InfoModel that will add the Attributes specified by the AttributeFactories to all LeifDataItems this InfoModel creates.

Figure 3A

<u>AttributeGetRequest</u>	The <u>AttributeGetRequest</u> class is used to package all the necessary parameters for getting the value an attribute.
<u>AttributeLockRequest</u>	The <u>AttributeLockRequest</u> class bundles attributes needed to gain access to locked <u>LeafDataItems</u> .
<u>Attributes</u>	The <u>Attributes</u> is a container for holding attributes.
<u>AttributeSetRequest</u>	The <u>AttributeSetRequest</u> class is used to package all the necessary parameters for setting an attribute.
<u>BaseDataItem</u>	The <u>BaseDataItem</u> is the first wrapper around raw data items.
<u>BaseInfoModel</u>	//PENDING(RK): Any method marked with "PENDING" in the JavaDoc will //likely be removed before XIS is released, in final form.
<u>BaseInfoModelServiceProvider</u>	The <u>BaseInfoModelServiceProvider</u> class will provide all the services available from the <u>BaseInfoModel</u> to the given <u>BeanContextServices</u> object.
<u>CollectionProperties</u>	This class populates <u>JAFMenus</u> for generic collections.
<u>DataItemActionManager</u>	The <u>DataItemActionManager</u> class provides some useful static methods for dealing with actions on <u>LeafDataItems</u> .
<u>DataItemActionManager.LeafReferenceActionListener</u>	This class is an <u>actionListener</u> to be used with <u>LeafReference</u> "Load" menus.
<u>DataItemMenuSet</u>	The <u>DataItemMenuSet</u> class is used by the <u>LeafAFUtilities</u> class to return essentially a <u>DataItem</u> popup menu with annotation.
<u>DataItemMenuSet.Entry</u>	The <u>DataItemMenuSet.Entry</u> class encapsulates a <u>DataItem</u> and it's menu, and also provides some convenience methods for adding additional menu items.
<u>DefaultWrapperAttribute</u>	The <u>DefaultWrapperAttribute</u> class is a generic attribute that is the superclass of all defaults in generated domain attributes.
<u>DisplayLabelAttribute</u>	The <u>DisplayLabelAttribute</u> class is used to display one or more <u>Attribute</u> values in conjunction with string literals specified by users.
<u>DisplayLabelData</u>	The <u>DisplayLabelData</u> class is used by the <u>DisplayLabelAttribute</u> to store a mapping of <u>LeafDataItems</u> to <u>DisplayLabelTemplates</u> .
<u>DisplayLabelTemplate</u>	The <u>DisplayLabelTemplate</u> class is used by the <u>DisplayLabelAttribute</u> to compute and store editing and rendering values for every <u>LeafDataItem</u> that has the attribute.
<u>DomainMethod</u>	Abstractly represents a <u>Domain Method</u> .
<u>DomainMethodDescriptor</u>	The <u>DomainMethodDescriptor</u> class is used to describe a <u>DomainMethod</u> .
<u>DomainMethodDescriptorFactory</u>	The <u>DomainMethodDescriptorFactory</u> class is a singleton class used to create or get <u>DomainMethodDescriptors</u> .
<u>DomainWrapper</u>	This class adds methods to <u>LeafDataItem</u> delegator that are useful in the domain wrappers.
<u>DynamicAttributes</u>	The <u>DynamicAttributes</u> class is used for storing dynamic attributes.

Figure 34c

<u>FieldMetaData</u>	FieldMetaData specifies sorting and subset criteria for an attribute.
<u>FieldMetaDatas</u>	The FieldMetaDatas class represents a collection of FieldMetaData for a data item.
<u>InfoModelDataItem</u>	The InfoModelDataItem allows views to wrap LeifDataItems and add/remove/modify attributes that will only affect that view.
<u>InfoModelSubset</u>	Typically when creating an InfoModel to nest within an existing InfoModel, which is done by ViewUIBeans, an InfoModelSubset is used.
<u>InvalidWrapperAttribute</u>	The InvalidWrapperAttribute class
<u>LeifDataItemComparator</u>	The LeifDataItemComparator compares LeifDataItems by AttributeDescriptor supplied by the user.
<u>LeifDataItemDelegator</u>	Implements the methods in LeifDataItem in a wrapper so you don't have to.
<u>LeifDataItemSorter</u>	The LeifDataItemSorter provides a default sorting tool for all LEIF LeifDataItem objects.
<u>LeifDataItemUpdate</u>	This classes is used with the LeifDataItemObserver.
<u>LeifInitialization</u>	The LeifInitialization class handles some standard initialization for most XIS Applications.
<u>LeifJAFUtilities</u>	The LeifJAFUtilities class provides some useful static methods for LEIF-related JavaBeans Activation Framework (JAF) processing.
<u>LeifJAFUtilities.LeifReferenceActionListener</u>	This class is an ActionListener to be used with LeifReference "Load" menus.
<u>LeifRequest</u>	The LeifRequest class is used to package all the necessary parameters for requesting information for a LeifDataItem.
<u>LeifTransaction</u>	The LeifTransaction class is used to construct a transaction.
<u>LockedLeifDataItem</u>	The LockedLeifDataItem class is used to enable locking on the data item.
<u>MethodRequest</u>	The MethodRequest class is used to package all the necessary parameters for invoking a DomainMethod for a LeifDataItem.
<u>MutableAttributeDescriptor</u>	Mutable subclass of AttributeDescriptor.
<u>ObserverSupport</u>	This class provides useful support for using the LeifDataItemObserver.
<u>RequestPool</u>	The RequestPool class is used to assist Object pooling.
<u>Resources</u>	The Resources class is automatically generated and must be public, but it is intended to be used only by Java's internationalization support classes.
<u>SelectableDataItem</u>	Creates a wrapper around a LeifDataItem for a SelectableInfoModel.
<u>SelectableInfoModel</u>	Manages selections for the selectable leif data items that are contained within this model.
<u>Translator</u>	A major design goal for XIS was to provide the ability to integrate existing data item classes without modifying them.

69/90

Figure 234

<u>TranslatorRegistry</u>	Provides a central location for maintaining Translators, extended Translators and locating Domain methods on data items.
<u>UndefinedAttribute</u>	The UndefinedAttribute class represents an undefined attributes.
<u>VisibilityAttribute</u>	The VisibilityAttribute class is an Attribute that is ready to use for LeiDataItem visibility.

Exception Summary	
<u>DataItemNotFoundException</u>	The DataItemNotFoundException class is an exception that can be thrown when trying to look up a data item from an id
<u>InvalidObjectSchemaException</u>	Signals that there was a problem with the creation or modification of an ObjectSchema.
<u>TranslatorException</u>	The TranslatorException class
<u>UnconvertibleAliasException</u>	Indicates that the requested attribute alias could not be calculated or converted.
<u>UndefinedLeafAttributeException</u>	Indicates that the requested attribute is not applicable for the object.
<u>UndefinedLeafMethodException</u>	The UndefinedLeafMethodException class indicates that the data item does not define the method.
<u>UnremovableAttributeException</u>	The UnremovableAttributeException class indicates an attempt to remove an Attribute that was defined by the raw data item (either by reflection or a Translator.) Only additional Attributes added to LeiDataItems can be removed.

Figure 35A

com.xls.leif.im

Interface InfoModel

All Superinterfaces:

BeanContextChildOwner, BeanContextChildOwnerDelegator, BeanContextProxy

All Known Implementing Classes:

InfoModelSubset

```
public interface InfoModel
extends BeanContextChildOwnerDelegator
```

The InfoModel interface is the interface that is used to convert raw data items into LeifDataItems. The InfoModel should hold each of these LeifDataItems created using weak references so that the data items can be cleaned up when they are no longer being used. //PENDING(RK): Any method marked with "PENDING" in the JavaDoc will likely be removed before LEIF is released in final form.

Since:

LEIF 4.0

Version:

\$Revision: 1.20 \$, \$Date: 2001/08/17 00:54:54 \$

Author:

David Almilli

Method Summary

void	<u>activateOneOfNService</u> (<u>Object</u> service) //PENDING(RK): This method will probably be removed from <u>InfoModel</u> ! Notify the <u>InfoModel</u> that the given service is the preferred service of its type, and that this particular object should be returned if its class is requested, until removed or until another object of the same type is passed to a future call to this method.
void	<u>addInfoModelListener</u> (<u>InfoModelListener</u> listener) Adds a listener to this <u>InfoModel</u> so that the listener will be informed of changes to the <u>InfoModel</u> .
void	<u>addOneOfNService</u> (<u>Object</u> service) //PENDING(RK): This method will probably be removed from <u>InfoModel</u> ! Add an object as a service to be retrieved by a call to <u>getService()</u> (via <u>BeanContext</u> APIs) on any class that this object implements or extends.
void	<u>clearSelection</u> () Clears the selection.
<u>LeifDataItem</u> []	<u>dump</u> () Gives a list of all the <u>LeifDataItems</u> currently in the <u>InfoModel</u> .
<u>EzContext</u>	<u>getEzContext</u> () Gets an <u>EZ Context</u> that corresponds to this <u>InfoModel</u> so the developer can use the <u>EZ APIs</u> .

Figure 35A

<code>LeifDataItem</code>	<code>getLeifDataItem(long uid)</code> This will attempt to lookup a <code>LeifDataItem</code> from an id.
<code>LeifDataItem</code>	<code>getLeifDataItem(Object rawDataItem)</code> This will wrap a raw java Object with a <code>LeifDataItem</code> wrapper so you can use it in leif as a data item.
<code>LeifDataItem</code>	<code>getLeifDataItem(Object rawDataItem, boolean create)</code> This will wrap a raw java Object with a <code>LeifDataItem</code> wrapper so you can use it in leif as a data item.
<code>LeifDataItems</code>	<code>getLeifDataItems(Object[] rawDataItems)</code> This convenience method will wrap an array of raw java Objects with <code>LeifDataItem</code> wrappers so you can use them in leif as <code>LeifDataItems</code> .
<code>InfoModel</code>	<code>getParentInfoModel()</code> Provides access to the parent <code>InfoModel</code> that this <code>InfoModel</code> delegates to.
<code>Object[]</code>	<code>getSelectedRawDataItems()</code> Gets the list of all the currently selected items for this <code>InfoModel</code> .
<code>Object</code>	<code>getSingleSelectedItem()</code> Get the selected raw data item, if only one.
<code>ViewHost</code>	<code>getViewHost()</code> Gets the <code>ViewHost</code> that this <code>InfoModel</code> is associated with.
<code>void</code>	<code>removeInfoModelListener(InfoModelListener listener)</code> Removes a listener from this <code>InfoModel</code> so that the listener will no longer be informed of changes to the <code>InfoModel</code> .
<code>void</code>	<code>removeOneOfService(Object service)</code> //PENDING(RK): This method will probably be removed from <code>InfoModel</code> ! Remove an object that was a service to be retrieved by a call to <code>getService()</code> (via <code>BeanContext</code> APIs) on any class that this object implements or extends.

Methods inherited from interface `com.xis.beans.beancontext.BeanContextChildOwnerDelegator`
`initializeBeanContextResources`, `releaseBeanContextResources`

Methods inherited from interface `com.xis.beans.beancontext.BeanContextChildOwner`
`getOwnedBeanContextChild`

Methods inherited from interface `java.beans.beancontext.BeanContextProxy`
`getBeanContextProxy`

Method Detail

`getLeifDataItem`

```
public LeifDataItem getLeifDataItem(long uid)
    throws DataItemNotFoundException
```

This will attempt to lookup a `LeifDataItem` from an id. If the UID is invalid or there isn't a `LeifDataItem` that already exists with that given UID, an exception will be thrown.

Parameters:

uid - the unique id for the raw data item.

Returns:

Figure 3

the LeifDataItem with the given UID

getLeifDataItem

```
public LeifDataItem getLeifDataItem(Object rawDataItem)
```

This will wrap a raw java Object with a LeifDataItem wrapper so you can use it in leif as a data item.

Parameters:

rawDataItem - the raw data that will be wrapped. (Note: this should not already be a LeifDataItem)

Returns:

the wrapped data item.

getLeifDataItem

```
public LeifDataItem getLeifDataItem(Object rawDataItem,
                                     boolean create)
```

This will wrap a raw java Object with a LeifDataItem wrapper so you can use it in leif as a data item.

Parameters:

rawDataItem - the raw data that will be wrapped. (Note: this should not already be a LeifDataItem)

create - if false and the LeifDataItem is not already in the model, don't create one and return null

Returns:

the wrapped data item, or null if "create" is false and not found

getLeifDataItems

```
public LeifDataItem[] getLeifDataItems(Object[] rawDataItems)
```

This convenience method will wrap an array of raw java Objects with LeifDataItem wrappers so you can use them in leif as LeifDataItems. Note that you can get an array of raw data items often from methods like getMembers(), so this is a useful method to have.

Parameters:

rawDataItems - the raw data objects that will be wrapped. (Note: the objects should not already be LeifDataItems)

Returns:

the corresponding wrapped data item array.

getEzContext

```
public EzContext getEzContext()
```

Gets an EZ Context that corresponds to this InfoModel so the developer can use the EZ APIs.

Returns:

the ez context for this info model

getSingleSelectedItem

```
public Object getSingleSelectedItem()
```

Figure 10

Get the selected raw data item, if only one. Else return null.

Returns:

the selected item if there is only one.

getParentInfoModel

public InfoModel getParentInfoModel()

Provides access to the parent InfoModel that this InfoModel delegates to. If there is no parent model then this will return null.

Returns:

the parent InfoModel

clearSelection

public void clearSelection()

Clears the selection.

getSelectedRawDataItems

public Object[] getSelectedRawDataItems()

Gets the list of all the currently selected items for this InfoModel

Returns:

all of the selected data items (as raw data items)

activateOneOfNService

public void activateOneOfNService(Object service)

//PENDING(RK): This method will probably be removed from InfoModel! Notify the InfoModel that the given service is the preferred service of its type, and that this particular object should be returned if its class is requested, until removed or until another object of the same type is passed to a future call to this method.

Parameters:

service - the object to become the preferred service

addOneOfNService

public void addOneOfNService(Object service)

//PENDING(RK): This method will probably be removed from InfoModel! Add an object as a service to be retrieved by a call to getService() (via BeanContext APIs) on any class that this object implements or extends.

Parameters:

service - the object to be returned when requested

Figure 3

removeOneOfNService

```
public void removeOneOfNService(Object service)
```

//PENDING(RK): This method will probably be removed from InfoModel! Remove an object that was a service to be retrieved by a call to getService() (via BeanContext APIs) on any class that this object implements or extends.

Parameters:

service - the object to be removed from service

getViewHost

```
public ViewHost getViewHost()
```

Gets the ViewHost that this InfoModel is associated with. If this InfoModel is not associated with a ViewHost then this will return null.

Returns:

the view host that is maintaining this InfoModel.

addInfoModelListener

```
public void addInfoModelListener(InfoModelListener listener)
```

Adds a listener to this InfoModel so that the listener will be informed of changes to the InfoModel.

Parameters:

listener - the listener to add

removeInfoModelListener

```
public void removeInfoModelListener(InfoModelListener listener)
```

Removes a listener from this InfoModel so that the listener will no longer be informed of changes to the InfoModel.

Parameters:

listener - the listener to remove

dump

```
public List<DataItem> dump()
```

Gives a list of all the DataItems currently in the InfoModel. It is highly recommended to use this method only if you absolutely have no other way of accomplishing the task you need to do. Please keep in mind that if you hold onto the DataItems contained in the array returned or if you hold onto the array itself, the items will not be removed from InfoModel until you release them. If you wish to hold onto them, you should wrap them in WeakReference objects.

Note: When you use the dump() method in combination with the addInfoModelListener so that you can keep track of the same set of DataItems as the InfoModel, you can synchronize on the InfoModel to get the dump and then add a listener to receive events of future changes.

Figure 6A

Example:

```
synchronized(infoModel) {  
    LeifDataItem[] dataItems = infoModel.dump();  
    infoModel.addInfoModelListener(this);  
    for (int i=0; i < dataItems.length; i++) {  
        processItem(dataItems[i]);  
    }  
}
```

Returns:

the list of LeifDataItems currently in the InfoModel.

See Also:

WeakReference

Figure 36A

Package com.xis.leif.event

This package contains classes for handling events in XIS.

See:

Description**Interface Summary**

<u>InfoModelListener</u>	The InfoModelListener is used to monitor changes to an InfoModel.
<u>LeifDataItemListener</u>	This class is used for listening to LeifDataItems for various events.

Class Summary

<u>AttributeChangedEvent</u>	An "AttributeChanged" event gets delivered whenever a data item changes an attribute value.
<u>ContainerAddedEvent</u>	A "ContainerAdded" event gets delivered whenever a data item is contained as a member in a new object.
<u>ContainerRemovedEvent</u>	A "ContainerRemoved" event gets delivered whenever a data item has been removed as a member from a containing object.
<u>DataItemReplacedEvent</u>	The DataItemReplacedEvent class is used to indicate member changes of a containing data item.
<u>InfoModelEvent</u>	The InfoModelEvent gets delivered whenever a LeifDataItem is created by the InfoModel, or when a LeifDataItem has been "lost" by the InfoModel.
<u>InfoModelEventSupport</u>	The InfoModelEventSupport support class provides basic support for managing listeners on an InfoModel.
<u>LeifDataItemAdapter</u>	The LeifDataItemAdapter class provides support for setting up a LeifDataItemListener on a data item.
<u>LeifEventSupport</u>	This is a utility class for XIS developers to use when they want to fire event changes.
<u>MemberAddedEvent</u>	The MemberAddedEvent class indicates that members were added to this data item.
<u>MemberEvent</u>	The MemberEvent class is used to indicate members changes of a containing data item.
<u>MemberRemovedEvent</u>	The MemberRemovedEvent class indicates that the members are being removed from the containing data item.
<u>ReferenceAddedEvent</u>	The ReferenceAddedEvent class indicates that LeifReferences were added to the LeifDataItem.
<u>ReferenceEvent</u>	The ReferenceEvent class indicates changes to the LeifReferences of the data item.

Figure 36B

<u>ReferenceRemovedEvent</u>	The ReferenceRemovedEvent class indicates that LeifReferences were removed from the LeifDataItem.
<u>ReferrerAddedEvent</u>	The ReferrerAddedEvent class indicates that a Referrer was added to the data item.
<u>ReferrerRemovedEvent</u>	The ReferrerRemovedEvent class indicates that a Referrer was removed from the data item.

Package com.xis.leif.event Description

This package contains classes for handling events in XIS.

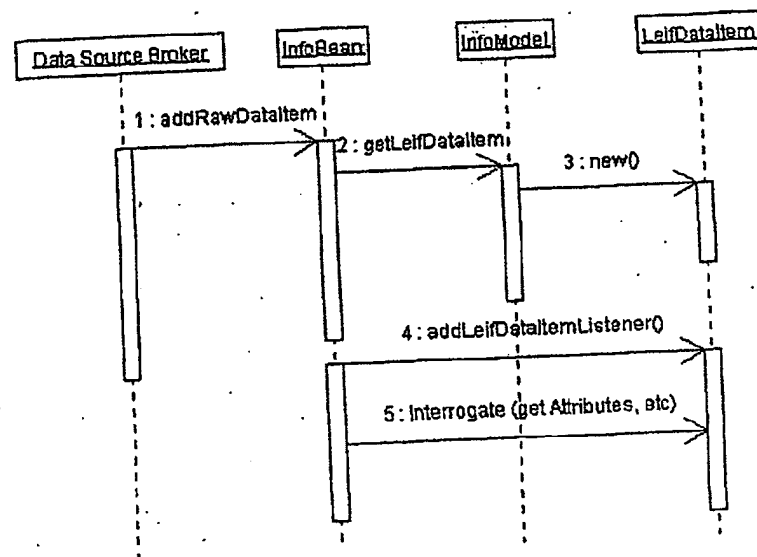


Figure 36C

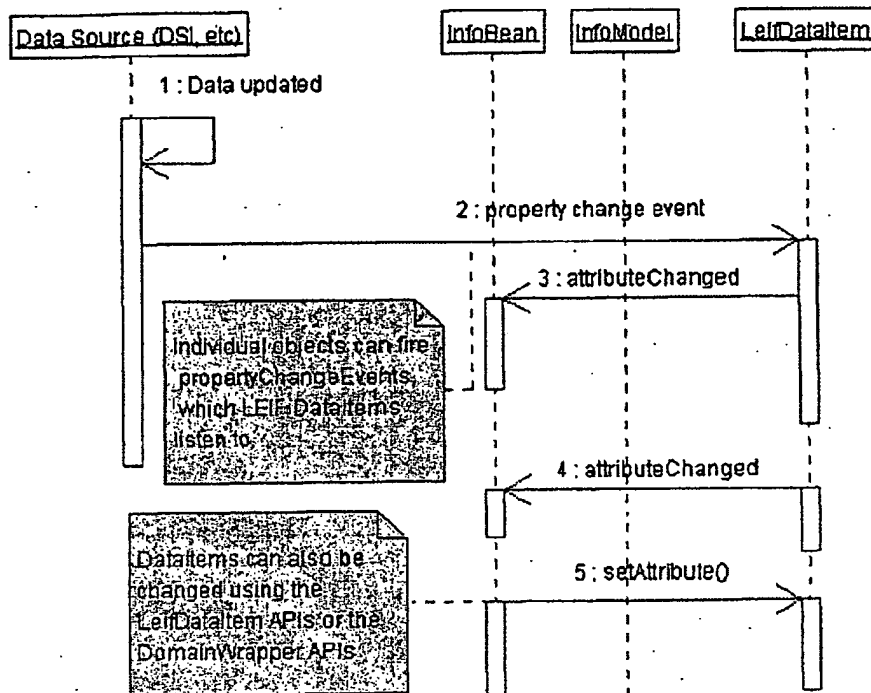


Figure 36D

Figure 37A

TestHarness.java

```
/* XIS Tutorial standalone sequence example 5 XIS interfacing. */

import javax.swing.JFrame;
import java.awt.event.WindowAdapter;
import java.awt.event.WindowEvent;
/*{*/
import java.awt.GridLayout;
import java.awt.Toolkit;
import java.awt.Dimension;
import java.awt.event.ComponentAdapter;
import java.awt.event.ComponentEvent;
import javax.swing.JSplitPane;
import javax.swing.JComponent;
/*}*/
import jclass.chart.JCChart;
import com.xis.leif.im.BaseInfoModel;
import com.xis.plot.PlotInfoBean;
import com.xis.plot.chartviews.LeifChartView;
/*{*/
import com.xis.table.TableInfoBean;
import com.xis.tree.TreeInfoBean;
/*}*/

public class TestHarness {

    public static void main(String[] args) {

        // the plugin manager is only required for more complex applications
        // involving multiple components integrated at runtime
        BaseInfoModel.setStartingPluginManager(false);

/*{*/
        HelloWorld hello1 = new HelloWorld("First HelloWorld object.");
        HelloWorld hello2 = new HelloWorld("Second HelloWorld object.");
        HelloWorld hello3 = new HelloWorld("Third HelloWorld object.");
        HelloWorld hello4 = new HelloWorld("Fourth HelloWorld object.");
        HelloWorld hello5 = new HelloWorld("Fifth HelloWorld object.");
    }
}
```

Figure 37B

Continuation of TestHarness.java

```

Object[] helloArray = new Object[] { hello1, hello2, hello3, hello4,
                                     hello5 };

// create table and plot infobean to display HelloWorld objects
TableInfoBean table = new TableInfoBean();
/*
PlotInfoBean plot = new PlotInfoBean();
plot.setChartType(JCChart.BAR);
// the alternatives are SCATTER_PLOT, PLOT, AREA, PIE, CANDLE,
// and STACKING_BAR, though not all will make sense in this example

// We can set the attribute for initial display on the plot;
// see step 3 for further comments.
plot.setYAxisAttribute(
    "com.xis.domains.movement.MovementDomain.speed");
plot.setDynamicAdjustment(true); // so axes track value magnitude
plot.setBarChartAdjusting(true); // needed in some cases for bar chart
*/

// a top-level frame as before to hold both plot and table side-by-side
JFrame tablePlotFrame = new JFrame("HelloWorld(s) Table/Plot");
// use a repaired JSplitPane (see below) to manage the two beans
SaneJSplitPane splitpane = new SaneJSplitPane(table, plot,
    new Dimension(table.getPreferredSize().width +
        plot.getPreferredSize().width,
        Math.min(table.getPreferredSize().height,
            plot.getPreferredSize().width)),
    0.50);
tablePlotFrame.getContentPane().add(splitpane);
tablePlotFrame.pack();
tablePlotFrame.setVisible(true);

// a tree infobean to display our HelloWorld objects
TreeInfoBean tree = new TreeInfoBean("HelloWorld(s) Tree");
tree.addRawDataItems(helloArray);

// a top-level frame to hold our tree infobean
JFrame treeFrame = new JFrame("HelloWorld(s) Tree");

```

Figure 37C

Continuation of TestHarness.java

```

// avoid placing the windows on top of one another if we can
int cutoffHeight = 424;
if (Toolkit.getDefaultToolkit().getScreenSize().getHeight() >
    cutoffHeight + 200) {
    treeFrame.setLocation(348,cutoffHeight+7);
}
/*}*/

// add a listener for window closing
treeFrame.addWindowListener(
    new WindowAdapter() {
        public void windowClosing(WindowEvent e) {
            System.exit(0);
        }
    }
);

// stick the tree infobean in the frame and display it
treeFrame.getContentPane().add(tree);
treeFrame.pack();
treeFrame.setVisible(true);

} // main

/*{*/
/*
 * This class overrides the default JSplitPane to provide a reasonable
 * resize behavior: maintain the left and right panels in the same
 * proportions.
 */
public static final class SaneJSplitPane extends JSplitPane {

    private int    lastWidth;
    private double lastDividerProp;

    public SaneJSplitPane(JComponent leftComponent,
        JComponent rightComponent,
        Dimension dims, double startProportion) {

```

Figure 37D

Continuation of TestHarness.java

```

super(JSplitPane.HORIZONTAL_SPLIT,
    leftComponent, rightComponent);
setSize(dims);

// Since the JSplitPane doesn't set the lastDividerLocation
// variable, nor does it provide any other easier way to maintain
// the split proportion on resize, we must track the divider
// location ourself.
lastWidth = dims.width;
lastDividerProp = startProportion;
setDividerLocation(startProportion);

// this listens for resize events on the splitpane and makes sure
// we keep same split proportions
addComponentListener(new ComponentAdapter() {
    public void componentResized(ComponentEvent event) {
        setDividerLocation(lastDividerProp);
        lastWidth = (int)event.getComponent().
            getSize().getWidth();
    }
});

// only way to know if divider moved by user is to listen for
// resize events on the components; this isn't foolproof (since
// resizes can come from other sources) but it works well enough
leftComponent.addComponentListener(new ComponentAdapter() {
    public void componentResized(ComponentEvent event) {
        // we add in getDividerSize() / 4 to compensate for a
        // bug in JSplitPane which doesn't take account of the
        // divider width in location-proportion conversions
        lastDividerProp = (double)(getDividerLocation() +
            (getDividerSize() / 4)) / lastWidth;
    }
});
}

}
/* */
}

```

Figure 38A

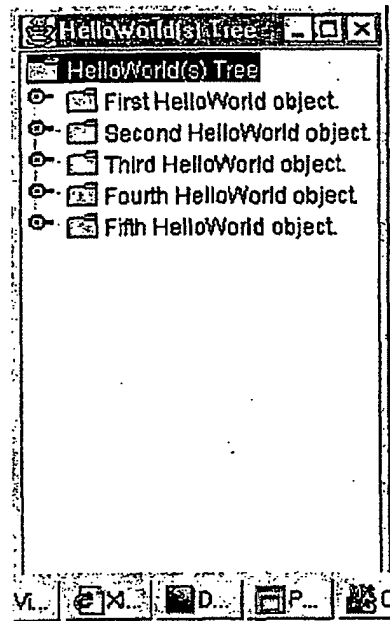



Figure 38B






Figure 38C

Property	Value
Name	First HelloWorld object
Speed (KTS)	1 ↑
Pen Color	□

Apply Reset

Figure 38D

Fig. 39A

com.xis.leif.im

Interface AttributeAlias

public interface AttributeAlias

The AttributeAlias indicates an alias from 1..n Attributes to a single Attribute, together with a precision level; the higher the precision, the better the alias. The alias allows a caller to query for one of the 1..n "from" attributes and get the value stored by the data item under the "to" attribute, possibly mediated by some conversion, such as a units transformation.

If the converted or calculated value cannot be determined, then the Attribute#getValue() method should throw an UnconvertibleAliasException. This Exception is a subclass of the UndefinedLeifAttributeException which is typically thrown by normal Attributes in this case, and it can provide a descriptive message indicating the source of the incompatibility.

The utility of attribute aliases may be seen by considering the following example:

The user has performed a query from an external data source and retrieved a set of Airfields, indexed by an ICAO identifier. The user now wants to get the list of Aircraft at one of the airfields. There is a local aircraft database with a foreign key field for Airfields, but that key is a WAC identifier, not ICAO.

Assumption: The application was NOT written ahead of time to know about these two databases or their ID types. Instead, what you have is an XIS "LeifDataItem" for the Airfield, and you have an XIS InfoBean for the Aircraft query form.

What you want to do is to copy (or "drop") the Airfield data item into the "WAC" field in the query form. In doing this, the Form will ask the data item for its "WAC" attribute (because this is all it knows about). It uses the "getWAC()" method from some domain (say, the AirfieldDomain).

The way this could work is that there would have to be an AttributeAlias defined to convert ICAO to WAC - or, more specifically, AviationDomain.ICAO to AirfieldDomain.WAC. The AttributeAlias returns an Attribute object that knows how to transform ICAOs to WACs (e.g., by accessing a conversion table). The Attribute, in turn, has a getValue() method to execute that transformation and return the WAC.

This process would be entirely transparent to the user, or even the caller, who would just see a result returned from the getWAC() method. In cases where the conversion was not possible, the UnconvertibleAliasException would be thrown, possibly providing informative information to the caller or user.

Finally, note that due to the way the mechanism is set up (using resources and a PluggableService), this AttributeAlias can be installed as a separate module without requiring any re-coding or re-compilation of the existing application.

Method Summary

<u>AttributeDescriptor</u> ()	<u>getAliasedFrom</u> () This indicates which AttributeDescriptors (which in turns means which Attributes) are required for the alias.
<u>AttributeDescriptor</u> ()	<u>getAliasedTo</u> () This indicates the AttributeDescriptor that this AttributeAlias is for.
<u>Attribute</u> ()	<u>getAttribute</u> () Get the Attribute object that is the alias Attribute.
<u>int</u> ()	<u>getPrecisionPriority</u> () This indicates the precision of the AttributeAlias.

Method Detail

getPrecisionPriority

```
public int getPrecisionPriority()
```

This indicates the precision of the AttributeAlias. The higher the number the better the alias. This number is used to determine which AttributeAlias to use when there are more than one alias for a given Attribute.

Returns:

the precision of the alias.

getAliasedFrom

```
public AttributeDescriptor[] getAliasedFrom()
```

This indicates which AttributeDescriptors (which in turns means which Attributes) are required for the alias.

Returns:

the list of descriptors required for this alias.

getAliasedTo

```
public AttributeDescriptor getAliasedTo()
```

This indicates the AttributeDescriptor that this AttributeAlias is for.

Returns:

the descriptor that this alias is for.

Fig. 39B

getAttribute

```
public Attribute getAttribute()
```

Get the Attribute object that is the alias Attribute. This attribute is responsible for performing the translation from the aliased from Attributes. The AttributeDescriptor of the Attribute **MUST** be the same AttributeDescriptor returned by getAliasedTo.

Returns:

the attribute that will do the translation.

Fig. 39C




Figure 40